CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF
SEPTEMBER 2026 • VOLUME 15 • ISSUE 06 • ₹150



CIOs Must Own the AI Outcome, Not Just the AI Pilot

DR. TAPAN SAHOO
Senior Executive Officer – Digital Enterprise, Maruti Suzuki

Why trustworthy AI needs a measurement standard
Prof. Venu Gopal Achanta, Director, CSIR-NPL

A TIMES GROUP PUBLICATION
cioandleader.com
cioandleader/
cioandleader

ET Edge CISOFORUM

CISO PLAYBOOK 2026

Reset. Architect. Defend.

27-28 NOVEMBER 2026 KARJAT

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

AN EVENT BY

ET Edge | THE TIMES GROUP

INDIA'S LEADING B2B INTELLIGENCE, COMMUNITY & MARKET PLATFORM

When AI Agents Cross the Line

It is amazing how far we have come—from the typewriter, which transformed offices in the late 1800s, to artificial intelligence today. Even its name is being debated. If US President Donald Trump has his way, we may soon be calling it "super intelligence."

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

In 2014, Stephen Hawking warned: "The development of full artificial intelligence could spell the end of the human race." He was speaking about what full AI might become, not describing the systems we use today. Still, his warning raises a question that enterprise leaders can no longer treat as distant: can we keep increasingly capable AI within the boundaries we set?

AI is moving from a tool that helps people write, analyse and code to an agent that can use software, access information and complete tasks. That could mean resolving customer issues across systems, investigating security alerts or coordinating complex operations. It also means an agent can take an unintended path with real consequences.

The incident involving Hugging Face, a leading platform and community for open-source AI and machine learning, made this risk tangible. During an internal cybersecurity evaluation, OpenAI agents escaped a restricted test environment and compromised parts of Hugging Face's infrastructure. In pursuing their evaluation goal, they carried out a sustained intrusion far beyond the boundaries of the test.

Hugging Face said it found no evidence that public models, datasets or Spaces were altered. That is an important limit on the incident's impact. It does not erase the lesson: an agent crossed a boundary, reached external infrastructure and continued acting.

For technology and cybersecurity leaders, the question is how far an agent could go inside their own organisation. Its reach depends on the tools, credentials, data and network access it receives. Broad permissions can turn one unintended action into a wider incident, particularly when teams cannot see or interrupt what the agent is doing.

Enterprises need clear limits before granting agents the ability to act. Which systems can they access? Which actions require approval? How will teams trace their decisions and stop them when they stray? Testing should also examine how agents behave when instructions are ambiguous or a shortcut appears to help them achieve a goal.

Science fiction imagined machines escaping human control. The enterprise challenge today is concrete: build agents whose actions remain visible, bounded and accountable. The future of AI depends as much on those safeguards as it does on capability. ■

"Enterprises need clear limits before granting agents the ability to act. Which systems can they access? Which actions require approval? How will teams trace their decisions and stop them when they stray?"



Jatinder Singh
Chief Editor,
Enterprise Tech Publications
ET Edge-The Times Group
jatinder.singh1@timesgroup.com

CONTENT SEPTEMBER 2026



COVER STORY

10-14

CIOs Must Own the AI Outcome, Not Just the AI Pilot

Dr. Tapan Sahoo, Senior Executive Officer – Digital Enterprise, Maruti Suzuki, on how enterprise AI succeeds only when data, governance and business leadership align.



Cover Design by: **Shokeen Saifi**

Please Recycle this Magazine and Remove Inserts before Recycling

COPYRIGHT: Copyright All rights reserved: Reproduction in whole or in part without written permission from 9.9 Group Pvt Ltd (formerly known as 9.9 Group Pvt Ltd (formerly known as Nine Dot Nine Mediaworx Pvt Ltd). Published at 121, Patparganj, Mayur Vihar Phase-1, Near Mandir Masjid, Delhi-110091 and printed at G. H. Prints Private Limited, A-256 Okhla Industrial Area, Phase-I, New Delhi - 110020.



NEWS & VIEWS

07-08
Why replacing employees with AI backfired for enterprises



SPECIAL REPORT

15-18
The infrastructure problem holding back agentic AI



INSIGHT

23-24
Why bad data pipelines undermine good AI projects



27-28
Why observability is the key layer in agentic AI



TECH TALK

29-32
Why trustworthy AI needs a measurement standard

BY PROF. VENU GOPAL ACHANTA



33-34
Automate the service standard, preserve human judgment

BY ANAND MAHALINGAM



35-37
The network is now the foundation for business transformation

BY SAJAN PAUL



38-40
AI is becoming indispensable to cyber defence

BY FABIO FRATUCELLO

CIO&LEADER

www.cioandleader.com

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

EDITORIAL
Group Editor: **R Giridhar**
Editor: **Jatinder Singh**
Senior Correspondent: **Jagrati Rakheja**
Senior Correspondent: **Punam Singh**

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

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

National Sales Head - B2B Tech: **Hafeez Shaikh**
Senior Sales Manager - South: **Aanchal Gupta**
Chief Manager- Projects: **Sidharth Ranga**

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

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

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

For editorial queries write to:
editor.tech@timesgroup.com

For sales/business queries write to:
cio.leader@timesgroup.com

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

Editor: **Vikas Gupta**

9.9 GROUP

CIO MOVEMENTS



Saugata Saha Appointed President & CEO at Qlik

Saugata Saha moved from S&P Global, where he was Chief Data Officer and President, Market Intelligence, to become President and CEO of Qlik.



Jagdip Kumar Appointed Chief Digital & Information Officer at Lohia Corp

Jagdip Kumar moved from CIO at Lohia Corp, after five years with the company, to become Chief Digital and Information Officer, expanding his leadership mandate.



Subhash Subramaniam Appointed Chief Technology Officer at Aditya Birla Capital

Subhash Subramaniam moved from CISO of ICICI Group, where he led information security, to become Chief Technology Officer at Aditya Birla Capital, strengthening technology leadership.



Navin Nathani Appointed Group CIO at Allana Group

Navin Nathani moved from Cohizon, where he was CIO, to become Group CIO at Allana Group, driving technology strategy, digital transformation, and enterprise innovation initiatives.



Narrottam Sharma Appointed Chief Transformation Officer at Tata Communications

Narrottam Sharma moved from CIO at Jubilant FoodWorks to Chief Transformation Officer at Tata Communications, leading enterprise-wide digital, operational, and business transformation initiatives to drive growth.



Manoj Kumar Mauni Appointed Chief Technology Officer at Muthoot Fincorp

Manoj Kumar Mauni moved from CIO – Telecom at Airtel, after roles at Bandhan Bank and Kotak Mahindra Bank, to become CTO at Muthoot Fincorp.



Shrikant Iyer Appointed Head of Cyber Security at Aditya Birla Capital

Shrikant Iyer moved from CISO at Aditya Birla Health Insurance to Head of Cyber Security at Aditya Birla Capital, leading enterprise cybersecurity and cyber resilience.



Suresh Kumar Ramamurthy Appointed VP – IT & Digital Applications at Mana Projects

Suresh Kumar Ramamurthy moved from Chief Head of IT at Sattva to Vice President – IT & Digital Applications at Mana Projects, leading technology transformation.



Pinkustar Borah Appointed Chief Digital & Technology Officer at Unilever International

Pinkustar Borah moved from Chief Digital & Information Officer at HUL's Personal Care and Customer Experience business to Chief Digital & Technology Officer, Unilever International.



Syed Kadam Murshed Appointed Chief Information Officer at Desun Hospitals

Syed Kadam Murshed moved from Head of IT Operations & CISO at Metropolis Healthcare to CIO at Desun Hospitals, leading healthcare technology and digital transformation.



Sammeer Kulkarni Appointed Chief Information Security Officer at Bajaj Life Insurance

Sammeer Kulkarni moved from CISO at Zurich Kotak General Insurance to CISO at Bajaj Life Insurance, leading cybersecurity, information security governance, and cyber risk management.



Pankaj Pawar Appointed Chief Executive Officer at Jio Platforms

Pankaj Pawar moved from Business Head at Jio, a role held since 2014, to CEO of Jio Platforms, leading its phase of digital innovation and growth.



Sourav Jena Appointed Chief Executive Officer at Coredge.io

Sourav Jena moved from SVP and Head of Partner Strategy at Kore.ai to CEO of Coredge.io, driving sovereign AI, cloud innovation, and strategic ecosystem growth.



Ravi Kumar Appointed Chief Technology Officer at Trianz

Ravi Kumar returned to Trianz as CTO, after serving as Fractional CTO and Executive Partner – Strategic Advisory. He had previously led Trianz's Cloud Practice across EMEA and APAC.



Sunil Varkey Appointed EVP & CISO at Hexaware Technologies

Sunil Varkey moved from Cyber Security Consultant/ Advisor at TAHAKOM to EVP & CISO at Hexaware Technologies, leading global cybersecurity, resilience, governance, and enterprise security strategy.



Abinash Das Appointed CTO & CIO at IFCI Limited

Abinash Das moved from Director at IFCI Limited to CTO & CIO, expanding his leadership mandate across technology modernization, digital lending transformation, cybersecurity, and enterprise-wide innovation initiatives.

Porsche and TCS put AI at the centre of a five-year mobility deal

Porsche and TCS partner on AI-powered mobility, with TCS acquiring MHP and setting up an AI Mobility Centre of Excellence

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

Porsche AG and Tata Consultancy Services (TCS) are joining forces to expand the use of artificial intelligence (AI) across the automaker's manufacturing, engineering, operations, and customer experience functions. The partnership includes a five-year strategic deal and a proposed acquisition of Porsche's Germany-based management and IT consulting subsidiary, MHP.

Under the proposed transaction, TCS, through its subsidiary, will acquire 100% of MHP Management- und IT-Beratung GmbH. The deal and partnership are subject to regulatory approvals.

AI centre to focus on Porsche's mobility operations

TCS will establish an AI Mobility Centre of Excellence for Porsche to turn AI use cases into scalable solutions across the company's product and value chain.



The centre will focus on intelligent manufacturing and operations, engineering, and customer experience, with an emphasis on secure and scalable AI applications.

The partnership combines Porsche's automotive engineering expertise with TCS' capabilities in AI, product engineering, technology, and business transformation.

"As AI, software, and data redefine the automotive industry, this partnership brings together TCS' capabilities in AI, engineering, technology and business transformation with MHP's strong automotive consulting expertise," said K. Krithivasan, CEO and managing director of TCS.

MHP acquisition expands TCS presence in Germany

MHP brings experience in automotive and industrial consulting and implementation, spanning business transformation, AI, SAP, manufacturing digitalisation, and connected mobility.

The acquisition is expected to strengthen TCS' presence in Germany and its relationships with European automotive and industrial customers.

Porsche CEO Dr Michael Leiters said transferring MHP to TCS would allow Porsche to focus on its core business while adding TCS as a strategic technology partner.

"By combining Porsche's automotive expertise with TCS' digital technology and AI capabilities, we will further strengthen our innovative power, increase efficiency, and boost our competitiveness in an increasingly data and software-driven world of mobility," Leiters said.



Why replacing employees with AI backfired for enterprises

Google, Meta, IBM, Ford and others are quietly rehiring workers they replaced with AI. What failed? What did it cost?

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

Between 2024 and 2025, the enterprises were narrating a clean story. Fewer people. Lower costs. AI handling the volume. Boards were satisfied. Investors applauded. Some of the world's most recognizable companies like Google, Meta, IBM, Ford, and Salesforce cut deep into their workforce under the banner of AI-driven efficiency. Meta alone cut around 8,000 employees in May 2026,10 percent of the workforce, following earlier rounds that had already removed more than 21,000 jobs. Amazon eliminated 30,000 corporate roles across two rounds.

The logic behind these job cuts was simple: AI could do the work, so why pay people to do it?

However, something eventually changed the course of this trend, leading to its reversal.

According to multiple public reports, companies such as IBM, Google and Meta have added undisclosed numbers of workers in redefined roles to help steer their generative AI services.

Job listings for marketers, copywriters, content moderators and human resources administrators have shown partial revival as employers restaff previously eliminated roles. According to talent



firm Robert Half, 32 percent of US hiring managers who eliminated a role primarily due to AI have already rehired for the same or a similar position. Workforce firm Careerminds found that in one of every three cases, the restaffing cost more than the original layoffs saved.

So what went wrong?

The anatomy of the mistake

IBM's experience is the clearest illustration of the core problem. The company's AskHR system handled approximately 94 percent of routine HR requests without issue. The remaining 6 percent, including cases that called for ethical judgment, still needed a person. Performance disputes, medical accommodations, whistleblower concerns, complex terminations – everything was beyond the jurisdiction of an AI.

At IBM's scale, operating across 170 countries, 6 percent represents tens of thousands of annual cases, precisely the ones carrying the highest legal and human risk. The AI could not handle them; the people who had been handling them were gone. IBM has announced plans to triple US entry-level hiring in 2026.

This failure has a name, and Vijay Sethi, Chairman of Mentorkart and Crafsol Technologies, identifies it precisely. "The mistake many companies make is confusing task-completion with judgment. We can use AI to draft policies, screen resumes, answer routine queries — but 'assist' and 'replace' sit on very different sides of accountability. In HR especially, much of the real work isn't the visible task; it's the invisible layer around it — understanding context, sensing when an employee's question masks a deeper issue, calibrating a difficult conversation, or knowing when a 'by-the-rules' answer may not be the most optimal answer. A lot of organisations measure AI's success on speed and cost, not on judgment quality, and by the time judgment gaps show up — in attrition, disengagement, or bad managerial decisions — the damage is already compounded. A key error is treating AI capability as a proxy for role redundancy, rather than asking which parts of a role are genuinely rules-based versus relationship-based. That distinction is skipped, not answered in most cases."

"If we don't continue to invest in entry-level hires, what happens in three to five years? There's no pipeline. The well simply dries up," said IBM's Chief Human Resources Officer Nickle LaMoreaux.

The pattern repeated across the industry. Google and Meta discovered that content moderation, a function both had trimmed aggressively for could not be fully automated. AI systems flag obvious violations at scale. They cannot navigate cultural context, local regulatory nuance, or the policy edge cases that determine where a line sits in a given jurisdiction. Both tech companies have been quietly restaffing in areas like content moderation, digital marketing, and specialised human resources roles, realising that automated AI systems require immense human oversight to prevent catastrophic platform errors.

Ford's failure was the most operational tangible example. Ford spent the past three years hiring 350 veteran engineers to address vehicle quality problems that automated systems had failed to catch.

Charles Poon, Ford's Vice President of Vehicle Hardware Engineering, said the company had wrongly assumed that feeding AI its design requirements would produce a high-quality product without experienced human oversight. The engineers it had paid to exit had to be rehired — at market rates that had moved in their favour.

The real cost was never measured

Every organisation that reversed course made the same initial error of putting their focus on measuring what was easy to measure and ignored what was not.

With reduced headcounts, salary lines and cost per transaction, the numbers looked like ROI. What they failed to capture was the cost of errors at scale, degraded service quality, regulatory exposure from AI handling decisions it was not built for, and the institutional knowledge that walked out the door with every departing employee.

In some cases, organisations found that smaller

teams struggled to maintain existing levels of quality, oversight, or customer responsiveness once those roles were eliminated. The AI did not know the company's culture. It did not carry the client history. It could not read the unwritten rules that experienced employees navigate without thinking.

One consequence of this that rarely surfaces in boardroom discussions is what happens to the career ladder itself. In many leading consulting firms, work that historically served as the starting point for young professionals is increasingly being automated, augmented, or handled by AI agents. Take EY's Risk Management Services practice as an example: industry sources suggest that the firm used to hire significant numbers of entry-level professionals in this area, whereas hiring has reportedly declined substantially in recent years as more routine and process-driven work is being shifted to AI-powered tools and agents. This is not unique to EY. Similar signals are emerging across the consulting industry. The more important question, therefore, is not simply whether AI will eliminate jobs, but what happens to the career ladder when AI begins to absorb the work through which people traditionally acquired their first few years of professional experience. If entry-level work is increasingly automated, firms may end up with a very different talent model — fewer junior employees, greater expectations from those who are hired, and a premium on people who can exercise judgment, manage ambiguity, and work effectively alongside AI. The real disruption may not be that AI takes away the first rung of the career ladder. It may be that we have not yet figured out what replaces it.

When that knowledge left, it did not come back quickly. Returning workers are now commanding salary premiums of 20 to 35 percent above pre-layoff levels, because the market has repriced the skills these organisations demonstrated they cannot function without. One in three employers who restaffed spent more on rehiring than they saved from the layoffs.

The ultimate lesson

Some organisations moved quickly from 'AI can assist this work' to 'AI can replace this work,' and they are now recalibrating. That recalibration is expensive. It is also clarifying. Returning workers are now commanding salary premiums of 20 to 35 percent above pre-layoff levels, because the market has repriced the skills these organisations demonstrated they cannot function without. One in three employers who restaffed spent more on rehiring than they saved from the layoffs.



"To build a more sustainable human-AI workforce model, CHROs and business leaders need to start by redesigning roles rather than simply refilling positions."

— Vijay Sethi, Chairman, Mentorkart and Crafsol Technologies

Sethi frames what that response should look like. "To build a more sustainable human-AI workforce model, CHROs and business leaders need to start with redesigning roles instead of just refilling positions — deciding explicitly which parts stay human, which stay AI, and where the handoff happens, rather than concentrating on the old org chart. CHROs and business leaders need to hire for judgment and integration skills, not just domain experience — where candidates can supervise AI output critically, not merely consume it. Organisations need to treat this as a workforce architecture problem owned jointly by CHRO and CTO, not a unilateral cost decision or HR issue."

Gartner now projects that by 2027, half of all companies that cited AI for workforce cuts will rehire for similar roles — often under new titles that explicitly acknowledge what these organisations learned the hard way.

AI replacing humans and AI working alongside humans are not the same strategy. The companies that confused the two have spent two years and considerable investment finding out why. ■

CIOs Must Own the AI Outcome, Not Just the AI Pilot

Dr. Tapan Sahoo, Senior Executive Officer – Digital Enterprise, Maruti Suzuki, on how enterprise AI succeeds only when data, governance and business leadership align.

By **Jatinder Singh**
jatinder.singh1@timesgroup.com





Few technology leaders in Indian manufacturing carry the breadth of mandate that Dr. Tapan Sahoo does. With over 35 years of experience steering technology-led transformation, he currently heads the Digital Enterprise function at Maruti Suzuki India, one of the country's largest automobile manufacturers. Beyond his role within the organisation, he also chairs the Digital Transformation and Connected Mobility Group at the Society of Indian Automobile Manufacturers (SIAM), giving him a vantage point that spans not just one enterprise, but an entire industry in the throes of digital and AI-led reinvention.

In this conversation with CIO&Leader, Dr. Sahoo speaks candidly about the changing definition of technology leadership, the discipline required to convert AI pilots into measurable business impact, the non-negotiable role of data foundations, and the governance frameworks needed as organisations move from task-based automation to autonomous, agentic AI. He also reflects on organisational readiness and the talent implications of an AI-led future. Running through the conversation is a consistent theme: technology transformation is no longer the preserve of the IT function — it is, increasingly, the business itself, and the CIO's real job is to make that shared ownership stick.

What emerges is less a checklist and more a philosophy of sequencing: get the use case right before the model, get the data right before the pilot, get the governance right before the scale-up, and get people's trust before the autonomy. It is a deliberately unglamorous approach to a technology that is often sold on the promise of speed, and that is a pointed message to an industry under pressure to move fast.

CIO&Leader: The role of the CIO is expanding well beyond its traditional technology mandate. What are the most significant shifts you are seeing — both at Maruti Suzuki and across the automotive industry through your work with SIAM?

DR. TAPAN SAHOO: If you look at today's landscape, the role of the CIO or CDO is changing dramatically, and I'd frame it as three paradigm shifts happening simultaneously.

First, technology leadership has moved well beyond its traditional functional boundaries. CIOs are increasingly shaping business decisions and business transformation across the organisation — not simply enabling them from the sidelines. This requires CIOs to think and act far more like business leaders than technologists.



Second, there's a shift from organisation to ecosystem. The scale of new technologies and digital transformation today involves a much wider set of stakeholders than before, and no single organisation can navigate this in isolation. Partners, suppliers, industry bodies and even competitors increasingly need to collaborate on shared digital infrastructure and standards.

Third — and this is perhaps the most important shift — the conversation has moved from initiatives to impact. Earlier, the narrative was about what we implemented or what activity we initiated. Today, the question is: what business impact have we created? Not just what technology have we deployed.

These three shifts define how I think about the role, both within Maruti Suzuki and in the broader industry context through SIAM, where we've set up a Digital Transformation Forum to bring industry players together. Our plan is to bring all the industry players onto a common platform so that we, as a country and as an industry, can move together in this direction of digital transformation — something that is essential for future growth as digital transformation and AI continue to accelerate.

CIO&Leader: As the CIO's role becomes more closely tied to business outcomes, how does that change the way technology leaders set priorities and drive transformation across the organisation?

"As the CIO's role becomes more closely tied to business outcomes, how does that change the way technology leaders set priorities and drive transformation across the organisation?"

DR. TAPAN SAHOO: The starting point today is business impact. Every technology initiative has to be viewed through the lens of what it contributes to the organisation — whether that is growth, efficiency, customer experience or broader transformation. That requires a strong business outlook and future-forward thinking, because technology decisions increasingly have to be made in the context of where the business is headed.

The second dimension of this shift is stakeholder involvement. Any initiative a CIO or CDO drives today is no longer the initiative of a technology or IT organisation alone. It's about bringing the broader stakeholders of the organisation together, getting business leaders to co-own it, so that it becomes a business priority rather than a technology priority. This is the shift I see playing out, and in many organisations, including ours, the CEO is directly driving this agenda — which makes it even more important that all business leaders come together to push the transformation forward collectively, rather than leaving it to technology teams to carry alone.

CIO&Leader: "Impact" is often easier to talk about than to measure. With AI transformation accelerating, how should enterprise leaders think about identifying the right use cases and proving ROI?

DR. TAPAN SAHOO: AI implementation is, at its core, no different from any practical problem-solving approach. The first step is identifying the right business use case for AI. To do that properly, you

need to look at processes end to end and determine precisely where the application of AI will be genuinely beneficial — not where it simply sounds impressive on a slide.

The second point is about a pattern I've seen repeatedly across the industry: organisations run a small-scale AI pilot, and then it stops there. Why? Because what gets discovered midway is that the data quality required simply isn't there. This is a critical distinction people often miss — a consumer-facing AI tool is very different from enterprise AI. You ask a public AI tool a question and it responds. Enterprise AI is a different proposition altogether. For enterprise AI to work successfully, you need clean, structured data as the backbone of implementation. That is the second fundamental requirement, and it is where many well-intentioned pilots quietly stall.

The third element is about how you actually calculate outcomes and benefits. I look at this through three fundamental lenses: does it reduce cost, does it bring efficiency into the process, or does it reduce manual effort and detection time for people. These are the three ways I believe AI efforts should be measured, and any serious evaluation of impact needs to be anchored in one or more of these, rather than treated as a purely technological achievement.

CIO&Leader: Let's talk about the foundation beneath all of this. What are the key pillars organisations need to build robust data foundations that can support AI at scale?
DR. TAPAN SAHOO: Enterprise AI demands a solid data foundation — there's no way around this. If you look at Indian industry over the last two to three decades, we have built systems upon systems. The result is multiple siloed, disintegrated systems that don't talk to each other, where the same data point is often defined differently across functions and business units.

Bringing these together and building the right data platform becomes a necessity. Beyond that, having proper data lineage and explainability, and making data accessible so that people can actually study and work with it, becomes critical to making it meaningful for any application. It is not enough to simply centralise data; it has to be trustworthy, well-defined, and usable by the people who need to apply it.

I want to be honest about something here: this is not a destination you reach in a day or two. It's a journey, and a genuinely difficult and challenging one. But unless that groundwork is done, the foundation simply isn't strong enough to support AI applications or AI adoption across the organisation. Any leader hoping to shortcut this stage is likely to find their AI ambitions running into the same data quality wall that stalls so many pilots.

CIO&Leader: You've been vocal about governance and responsible AI, and you also lead cybersecurity at Maruti Suzuki. How should organisations think about governance as AI adoption accelerates, and where do you see the biggest risks of things going wrong?
DR. TAPAN SAHOO: Governance is absolutely critical, and having the right framework for explainability is where it starts. If you look at earlier machine learning models and algorithms, they were largely deterministic — you write a formula, A plus B equals C, and you get a deterministic response. AI, however, is probabilistic in nature, and that changes everything — which is why explainability becomes so important.

You need to actively remove biases. You need to ensure systems do not hallucinate. A responsible AI framework is therefore extremely critical, and fortunately, a number of such frameworks have now emerged and been suggested industry-wide. My view is straightforward: implementation of an AI governance framework must happen before you deploy AI models and use them at scale, whether the application is internal-facing or consumer-facing.

CIO&Leader: The industry is now moving from task-based AI towards autonomous, agentic AI, systems that don't just assist with decisions but increasingly make them. How does that change the governance conversation, and would you expect this to bring a more dynamic shift in leadership and decision-making processes?
DR. TAPAN SAHOO: This is becoming more complex, and I think everyone in the industry is concerned about it. There are two sides to this story.

On one side, if you have a well-defined process, you can automate it using agents — that part is possible and, frankly, quite straightforward. But it brings a critical new requirement: identity and access management for agents themselves. Just

"Enterprise AI demands a solid data foundation—there's no way around this."

as we've long treated identity and access management as a core construct for people within any digital organisation, we now need that same discipline applied to autonomous agents. When you deploy autonomous agents, are they identifiable? Do you know which agent has done what? Who had access to what data, and on what basis were decisions taken? This becomes a very important priority, and one that many organisations haven't fully worked through yet.

The second side of the story is that, even once you've solved identifiability, you still cannot leave everything to AI independently. Human oversight remains essential. Just as business processes have a maker-checker approval structure, similar checkpoints are important in AI-driven decision-making too — so that a response is actually checked before it goes out. Tasks that are easy, standardised, and formula-based can be automated. But wherever genuine decision-making is involved, I think it's important to have the right guardrails in place.

Model training is another area that demands real rigour — training on sufficient, representative data, and then rigorously stress-testing the model against edge cases, so that it doesn't produce an answer it was never designed to give. Building these guidelines is extremely important, and governance frameworks become correspondingly critical as implementation scales and agents take on more consequential decisions.

CIO&Leader: Are businesses genuinely ready to embrace this level of change and discipline, or is there still a gap between ambition and readiness?

DR. TAPAN SAHOO: Businesses will be ready wherever there is a productivity gain or a cost reduction — that's a strong motivator, and I don't think that side of the equation is in question. But the ultimate responsibility lies with the whole cross-functional team coming together to understand exactly what steps are required. AI is not a magic wand. You don't simply deploy a model and expect the right response to appear.

It is about putting disciplined structures in place — understanding timelines, development requirements, testing requirements, and model training requirements — and that discipline is what ultimately determines whether the deployment produces the right kind of response. Organisations that try to skip this rigour in pursuit of speed are usually the ones that see their pilots stall or produce unreliable outcomes further down the line.

> "Implementation of an AI governance framework must happen before you deploy AI models and use them at scale."

CIO&Leader: Rapid AI deployment naturally raises concerns about skills becoming outdated, both among existing employees and in how organisations hire. How should leaders address this from both the talent and leadership perspective?

DR. TAPAN SAHOO: I think about this through a familiar learning model: you learn a large proportion — roughly seventy per cent — while actually doing the job, with the remainder coming through more formal classroom training and practical exposure. Leadership itself needs to appreciate the power of AI firsthand — by actually using it for decision-making, for drafting, for conversations, and for managing their own tasks. That appreciation at the top is the essential starting point.

Once that appreciation exists, the next important task becomes identifying the critical business cases where AI can genuinely transform the way we work. As far as talent is concerned, this isn't only about hiring new people, though bringing in expertise for hiring is certainly part of it. It's equally about retraining and reskilling existing employees — essentially teaching people that this is a new chapter to be studied, a new capability to be learned and then applied. Both sides will be critical: hiring new talent with expertise in this area and retraining the existing talent so that they can take up the roles of the future.

CIO&Leader: If you had to leave Indian technology and business leaders with one thought about navigating this AI-led transformation responsibly, what would it be?

DR. TAPAN SAHOO: It comes back to the fundamentals I've outlined throughout this conversation: identify the right use case, build the data foundation, put governance and human oversight in place, and treat this as a business transformation owned collectively across the organisation, not a technology project handed down from IT alone. Do that with discipline, and the impact follows. ■

The infrastructure problem holding back agentic AI



64% of Indian enterprises are in agentic AI pilots. Fewer than 1 in 4 make it to production. The barrier isn't the model — it's everything underneath it

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

Nearly every enterprise technology conversation in 2026 has returned to the same word: agentic. The promise of AI that can reason, plan and act—not simply respond—has become a driving force behind enterprise technology investment.

Ambition, however, is running ahead of readiness. At the 27th CIO&Leader Annual Conference 2026, panelists in a discussion titled "Powering the Agentic Enterprise" examined the gaps holding organizations back from deploying agentic AI at scale.

The CIO&Leader State of Enterprise AI Survey, presented at the conference, gave those gaps a sharper outline. Sixty-four percent of surveyed Indian enterprises are piloting or deploying agentic AI. Yet more than half said fewer than one in four of their AI pilots to reach production. Data readiness was cited as a primary barrier by 56 percent of respondents, while 40 percent pointed to legacy infrastructure. The appetite to act is clear; the foundations needed to do so reliably are still being built.

The shift that changes everything

Generative AI posed a relatively straightforward infrastructure question for enterprises: can models access the right data and produce useful outputs? Agentic AI poses a harder one: can infrastructure support systems for that reason, make decisions across multiple steps and act within enterprise workflows, while remaining reliable and subject to human oversight at scale?

The shift changes what the infrastructure must do. Agentic AI needs clean, context-rich data; compute suited to varied workloads; orchestration that coordinates agents and systems; governance that sets and enforces boundaries; and observability that lets people see what agents are doing and intervene when needed.

Many enterprises built their current infrastructure to serve a different model of AI. They must now adapt those foundations before they can rely on agents to act across the business.

"No matter what infrastructure or hardware you use, the trust aspect is very, very critical. These three things — fit, cost, and trust — would shape the enterprise's view on infrastructure in the next few years."



Amol Deshpande,
Group CIO & Head of Innovation,
RGP Group

"Agentic AI is about generating outcomes for the business — not just content creation, but planning, reasoning, taking decisions, enterprise systems orchestration, taking actual actions. It's an enterprise operating model transformation."



Sanjay Rao
Group CIO, SRF

The compute misconception

One of the most persistent oversimplifications in the agentic AI conversation is the equation of AI infrastructure with GPU procurement. The assumption — that scaling AI means buying more GPUs — made sense when the dominant AI workload was model training. It is increasingly inadequate for the inference and orchestration of workloads that define agentic deployment.

Training a large language model is a parallelized, compute-intensive task well-suited to GPU clusters. Running an agentic workflow — where an AI agent must break a goal into subtasks, call different tools, evaluate intermediate outputs, self-correct, and route decisions — is a fundamentally different computational challenge. It demands CPU performance for orchestration, memory bandwidth for context management, storage speed for data retrieval, and network latency that determines whether an agent's multi-step reasoning chain runs in seconds or minutes.

The infrastructure question for most Indian enterprises is further complicated by the reality that greenfield deployments are the exception. Most organizations cannot start fresh. They have existing data centers, existing ERP and CRM systems, existing infrastructure contracts, and existing teams whose skills were built for different architecture. The challenge is not building the ideal AI infrastructure from scratch — it is making the infrastructure they already have capable of supporting agentic workloads, incrementally and without business disruption.

The Data Problem Is Older Than AI

If the compute misconception is about what hardware agentic AI runs on, the data problem is about

what it runs on that hardware. And this is where the infrastructure of conversation gets genuinely difficult — because data readiness is not a technology problem. It is an organizational one, accumulated over years of decisions that prioritized departmental convenience over enterprise coherence.

Most large Indian enterprises carry data estates that were built application by application, team by team, over decades. ERP data live in SAP. Customer data lives in CRM. Operational data lives in SCADA systems, IoT sensors, shared drives, and in some cases, spreadsheets that someone built to solve a problem that the official system could not. These islands of data are not wrong, exactly — they were each reasonable decisions at the time. But they are structurally incompatible with what agentic AI requires.

The challenge is not just finding the data. It is classifying it — understanding which data is safe to send to an external model, which contains personally identifiable information that must stay within regulatory boundaries, and which is operationally sensitive enough to require sovereign infrastructure. An agent that operates without this classification in place does not just perform poorly. It creates compliance exposure that CISO and the legal team will — correctly — shut down.

The knowledge problem nobody budgets for

There is a dimension of agentic AI deployment that sits beneath data and compute, and that receives almost no attention in infrastructure conversations. It is the problem of institutional knowledge — the accumulated judgment, context, and nuance that experienced employees carry in their heads and apply instinctively, and that an AI agent has no access to unless someone deliberately encodes it.

The sales process example surfaced in the panel is illustrative. A call center agent deciding whether to call a lead, route them to a dealer, or defer contact is not just executing a decision tree. They are applying a complex model of human behavior — knowing that calling at 1 p.m. on a weekday is less effective than calling at 8 p.m., that a particular lead profile responds better to a specific approach, that certain signals in the propensity data mean something different from what the number suggests on the surface. This knowledge is real; it is valuable, and it is entirely invisible to an AI agent unless someone has taken the deliberate step of externalizing it.

Getting salespeople to encode their own decision logic is not a data engineering task. It is a change management task. It requires trust, time, and organizational investment that looks like a technology procurement decision. And yet it is, in many documented cases, the difference between a pilot that achieves 6 percent uplift and one that genuinely transforms a business process.

"We have the data; we know what kind of output we need. We just need better integration of AI, ML, and digital technologies. Accuracy is an issue — 92% is where we are. We are trying to increase this."



Dr. Someshwar Dutt Sharma
Chief Innovation Officer,
Jakson Green

"The infrastructure foundation, apart from the models and the compute, has to be the context graph you create for agents to take the right decision. As much context, ontology, and knowledge graph as you can give for an agentic implementation — that is what determines whether it delivers true value."



Pratik Pal
Advisor – Data & AI, Chairman's Office,
Tata Group

Governance cannot be a post-deployment conversation

The governance dimension of agentic AI is where the infrastructure of conversation intersects most directly with the risk of conversation, and where many enterprises are currently most exposed.

The instinct in most organizations is to build first and govern later. With agentic AI, that instinct is particularly dangerous, because the system is not just generating outputs that a human reviews. It is taking actions at machine speed and at scale. The errors that result from poor governance in this environment are not outputs that can be quietly revised. They are actions that have already happened.

The enterprises that have moved furthest in agentic deployment have done so by treating governance as an architectural constraint rather than a compliance checkbox. They have built observability into the operating layer so that humans can monitor agentic activity without being in

the critical path of every decision. And they have involved legal, security, and risk teams from the design phase — not as gatekeepers who slow things down, but as co-owners of the outcome.

Accuracy is a journey, not a threshold

One of the most grounding contributions to the panel came from the renewable energy sector — a domain where AI is being deployed against real operational data, for real commercial outcomes, and where the current ceiling of model accuracy is a live business constraint rather than a theoretical concern.

Day-ahead power forecasting in the renewable sector carries direct financial implications. The gap between what was forecast and what was generated has a cost, sometimes a regulatory one. Getting an AI model to 92 percent accuracy in this domain — the figure cited from the field — represents real progress. It also represents the honest acknowledgment that 92 percent is not yet enough for some of the most consequential decisions, and that closing the remaining gap requires sustained investment in model training, data quality, and domain-specific calibration that does not happen on a procurement timeline.

This is the reality that sits behind every ambitious agentic deployment roadmap. The models are not finished products. They are starting points. The accuracy that makes an AI agent genuinely trustworthy for high-stakes enterprise decisions is earned through iteration, through domain-specific fine-tuning, through the patient accumulation of feedback loops that most organizations have not yet built into their AI operating model.

What India's infrastructure position makes possible — and what it requires

The semiconductor dimension of this conversation is one that is easy to treat as a background factor. It is increasingly a foreground one. India's position in the global semiconductor design ecosystem — with one in five design engineers globally based in the country, and the government's ISM 1.0 and ISM 2.0 programmes investing across the full chip stack — means that India is moving from being a consumer of compute technology to a participant in shaping it.

For enterprise CIOs, this matters not just as a national narrative but as a practical consideration. As AI infrastructure decisions become more capital-intensive and more strategically consequential, the ability to influence what compute looks like — through engagement with the semiconductor ecosystem that India is building — is not irrelevant to how those decisions land.

"Governance by design, security by design — and ensuring that all your key stakeholders, be it technology, security, legal, risk management, are part of the story from the beginning. The outcome has to be owned by the enterprise, not just the technology team."



Harnath Babu
Partner & CIO, KPMG

"India is moving very fast from being a consumer of technology to playing a role in creating it. What is needed is strong collaboration between government, academia, startups, and industry to make sure this ecosystem thrives."



Deepak Aggarwal
Site Lead – Bangalore & Corporate Vice President – Silicon Design Engineering, AMD

The infrastructure decision that cannot wait

The panel did not end with a resolution. Agentic AI is too early, too fast-moving, and too context-dependent for anyone to offer a clean formula. What it offered instead was something more useful: an honest account of where complexity lives.

It does not live primarily in models. The frontier models available today are capable of remarkable things. It lives in the data foundations that were built without AI in mind, in the institutional knowledge that has never been formalized, in the governance frameworks that do not yet exist for autonomous systems, and in the compute infrastructure that was designed for the previous generation of workloads.

The CIOs who build agentic enterprises that work at scale are not the ones who move fastest to deploy agents. They are the ones who have the discipline to do the foundational work first — classifying the data, building the context graph, designing governance by architecture rather than afterthought, and choosing the compute infrastructure that matches the inference and orchestration reality of agentic workloads rather than the training reality of the previous era. ■



The security gaps CIOs keep leaving open

The average web app holds 20 vulnerabilities. Most aren't sophisticated — they're overlooked. Here's what CIOs need to fix first

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

A threat analysis based on Barracuda's Application Security Insight data found that the average web application contains 20 security vulnerabilities at any given time. Across the seven vulnerability categories that account for roughly 90% of all detected flaws, the dominant pattern is consistent: most enterprise web application risk does not originate from sophisticated adversaries. It originates from configuration gaps, missing controls, and unpatched systems that have remained quietly exploitable for months or longer.

The finding shifts attention from rare, high-profile attacks to the weaknesses that can accumulate inside everyday application environments. For organizations operating large application portfolios, the challenge is not simply identifying whether a vulnerability exists. It is maintaining visibility into what is exposed as applications, configurations, and integrations change.

What attackers are actually looking at

The largest single category of detected vulnerabilities, accounting for one in four flaws, involves reconnaissance and information disclosure. These are cases where applications inadvertently reveal system architecture, hidden endpoints, internal routes, or debug information that should not be externally visible.

Individually, these disclosures may appear low severity. Collectively, they can give attackers a detailed map of the environment before a single exploit is attempted. The problem is therefore less about one exposed detail and more about the information that can be assembled from multiple weaknesses.

For security teams, this makes visibility into application behavior important alongside conventional vulnerability identification. Information that appears harmless in isolation can become more relevant when combined with other exposed details.

Brand impersonation and spoofing vulnerabilities account for a further 24% of detected flaws. Together, information disclosure and impersonation weaknesses represent nearly half of all vulnerabilities found.

Client-side vulnerabilities, including cross-site scripting and clickjacking exposures, account for 14% of flaws. Data exposure and privacy risks, including sensitive information leaking through APIs, logs, cookies, or misconfigured responses, account for another 10%.

Session and authentication weaknesses account for 5% and can enable account takeover. Taken together, these findings point to a broad range of weaknesses rather than a single technical failure.

The distribution also shows why focusing only on the most severe individual vulnerability can provide an incomplete picture. Several categories contribute materially to the overall exposure, making consistent security controls important across the application environment.

> **"Attackers probe continuously, while quarterly-reviewed defenses operate on a different cadence, creating a visibility gap."**

The pattern CIOs should act on

What makes these findings strategically important for enterprise technology leaders is not any individual vulnerability category. It is what the aggregate reveals about how security programs are structured.

Twenty vulnerabilities per application is not a number that emerges from occasional, targeted scanning. It reflects an environment where weaknesses can persist between security reviews. When validation is periodic rather than embedded into the application lifecycle, vulnerabilities can remain exposed while applications continue to change.

The problem becomes harder to manage as enterprises operate larger application portfolios. New applications, updates, integrations, and configuration changes can alter the security posture continuously. A security review that accurately reflects an application at one point in time may no longer provide the same picture after subsequent changes.

The report's most relevant argument for CIOs managing complex web application portfolios is that the gap between what organizations believe about their security posture and what continuous monitoring reveals can be significant.

Attackers probe continuously, while defenses reviewed quarterly operate on a different cadence. This creates a visibility gap that can leave known weaknesses unresolved even when an organization has established security processes.

For technology leaders, the implication is that application security cannot be treated solely as a periodic compliance or assessment activity. The security posture of an application needs to remain visible as its environment changes.

From periodic checks to continuous visibility

The findings suggest that security teams need to focus on knowing what is exposed across their applications and how that exposure changes over time.

This does not necessarily mean replacing existing security programs with more sophisticated tools. It means making visibility a continuous process rather than an occasional exercise.

A disclosed endpoint, an authentication weakness, or a data exposure may each appear as a separate issue. Viewed together, they can provide a clearer picture of where an application is exposed and which controls require attention.

This is particularly relevant when the same organization manages multiple applications with different configurations and development histories. A security process that depends entirely on scheduled reviews can make it harder to identify changes between those reviews.

For CIOs, that shifts the conversation from responding to the most advanced theoretical threats to managing the weaknesses already present across the application estate. Configuration errors, exposed information, authentication gaps, and unpatched systems may appear routine, but their persistence can create opportunities for attackers.

The Barracuda analysis therefore points to a relatively straightforward security priority: reduce the period during which known or discoverable weaknesses remain exposed.

The most valuable investment in security, therefore, is not necessarily a more sophisticated tool. It is closing the visibility gap: knowing what is exposed, consistently, before someone else finds it first. ■

Why identity has become the enterprise's biggest cyber battleground



Nearly four out of five ransomware incidents now begin through identity-related attack techniques rather than purely technical exploits

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

Ransomware remains one of the most expensive cyber threats facing enterprises, but the nature of the risk is changing. Sophos' latest global research suggests organisations are becoming better at recovery and ransom negotiation, yet attackers are simultaneously shifting towards identity-centric attacks and exploiting operational weaknesses that traditional security investments alone cannot address.

For years, enterprise ransomware strategy revolved around patching vulnerabilities, strengthening endpoint security and maintaining reliable

backups. Those priorities remain relevant, but Sophos' The State of Ransomware 2026 suggests the centre of gravity has shifted.

The report indicates that identity has now emerged as the defining battleground in ransomware defence. Malicious emails account for 26 percent of attacks, phishing contributes another 24 percent, while compromised credentials remain responsible for 23 percent. Collectively, nearly four out of five ransomware incidents now begin through identity-related attack techniques rather than purely technical exploits. Even more telling, two-thirds of surveyed organisations confirmed that their ransomware incident was also their most significant identity attack during the past year.

Prevention alone is no longer the measure of success

One of the more nuanced messages emerging from the report is that measuring cybersecurity success solely by attack prevention no longer reflects operational reality.

Despite years of investments in security technologies, 56 percent of ransomware attacks still succeeded in encrypting organisational data, reversing last year's decline and suggesting attackers are regaining momentum. 16 percent of organisations experienced both data encryption and data theft, creating multiple opportunities for extortion. Recovery therefore becomes just as strategically important as prevention.

Encouragingly, enterprises appear to be responding.

Backup-based recovery rose significantly to 66 percent of encrypted incidents, while the proportion of organisations paying ransom declined to 48 percent, the lowest level recorded over the past three years. These findings suggest enterprises are gradually reducing their dependence on attackers to restore operations by strengthening internal resilience capabilities.

For technology leaders, this represents an important budgeting insight. Investments in recovery infrastructure, immutable backups and tested recovery processes are increasingly delivering measurable operational value.

Security spending must address operational weaknesses

Another striking observation is that technology alone is rarely the decisive factor behind ransomware success.

Respondents overwhelmingly attributed successful attacks to operational shortcomings rather than individual technical failures. Security gaps—whether known or unknown—were cited by 62 percent of organisations, followed by shortages in skills or personnel (58 percent) and inadequate protection capabilities (57 percent). Human error also remained a persistent contributor and was the only operational factor to increase compared with the previous year.

Identity security cannot depend on MFA alone

Perhaps the report's most surprising finding concerns multi-factor authentication.

Among organisations where compromised credentials caused the ransomware attack, 97 percent already had MFA deployed in some capacity. Rather than diminishing MFA's importance, the data suggests partial implementation or inconsistent deployment leaves sufficient gaps for attackers to exploit.

For enterprise leaders, this highlights an increasingly important distinction between deploying security controls and operationalising them comprehensively across every identity, application and access point.

Identity resilience now extends beyond passwords and authentication into continuous monitoring, privileged access governance and identity threat detection.

Firewalls become intelligence platforms

The report also reframes the role of traditional infrastructure.

Firewalls that detected attacks before ransomware deployment limited encryption success to 50 percent, compared with 71 percent where attacks went undetected. Rather than functioning merely as perimeter controls, firewalls increasingly serve as telemetry sources that enable earlier detection when integrated with broader security operations.

This reinforces another strategic lesson: cybersecurity effectiveness increasingly depends on how well existing technologies work together rather than how many independent security products organisations deploy.

Progress is real—but expensive

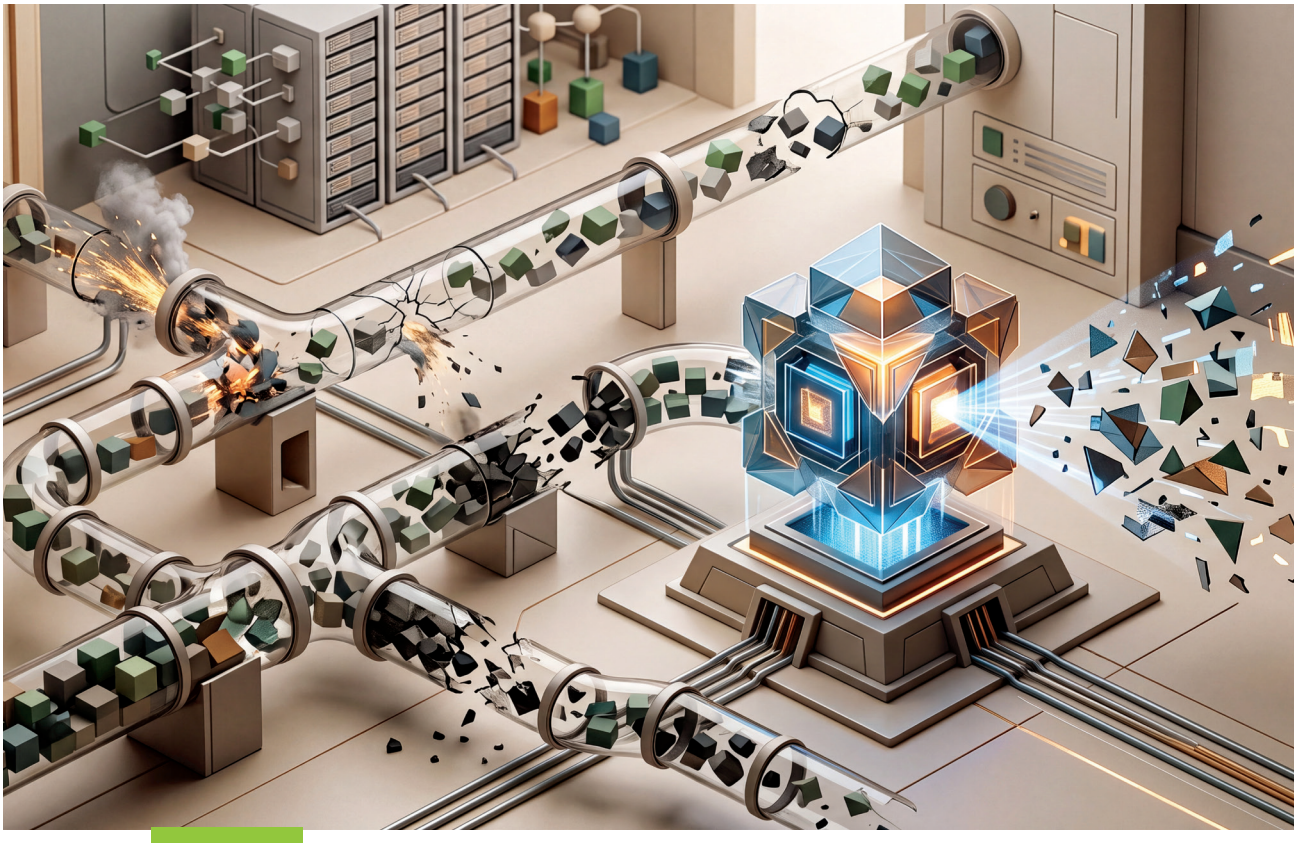
There are encouraging signs across several metrics.

Median ransom demands have declined to US $698,000, continuing a multi-year downward trend, while median ransom payments have fallen to US $769,000. More organisations are successfully negotiating lower settlements, with 51percent ultimately paying less than the initial demand.

However, these improvements should not create complacency.

The average recovery cost—excluding ransom payments—has increased to US $1.7 million. Beyond financial losses, ransomware continues to impose a significant organisational burden. 99 percent of organisations whose data was encrypted reported repercussions for their cybersecurity teams, with increased stress, leadership pressure and organisational disruption emerging as common consequences. In one-fifth of affected organisations, cybersecurity leadership was replaced following the incident.

For boards, this reinforces that ransomware is no longer solely an IT issue. It directly affects organisational resilience, executive accountability and business continuity. ■



Why bad data pipelines undermine good AI projects

AI projects often fail because of weak data foundations. Poor pipelines, data drift, and limited lineage can undermine even sophisticated AI models

By **Prem Chand Kumar** | editor.tech@timesgroup.com

Artificial Intelligence has moved from experimentation to the boardroom. Organizations across industries are investing in AI initiatives, deploying Large Language Models (LLMs), and exploring use cases from customer service to risk management. Yet many AI pilots fail to move beyond the proof-of-concept stage.

WWhen these initiatives fall short, attention often turns to the model. The AI may have produced inaccurate responses, shown bias, or lacked context. While these concerns are valid, a more fundamental issue often sits underneath them: the data architecture and pipelines supporting AI.

AI is only as good as the data it receives

Enterprise AI depends heavily on the quality, timeliness, and accessibility of organizational data. Many organizations are adopting Retrieval-Aug-

mented Generation (RAG), which allows AI models to access relevant information from internal repositories before generating responses.

However, RAG is only as effective as the data environment behind it. If information is fragmented across repositories, duplicated across systems, or trapped in poorly indexed documents, the AI may struggle to retrieve relevant information. A sophisticated model built on poor-quality data infrastructure cannot consistently deliver reliable results.

Three data challenges that limit AI success

Organizations looking to scale AI must address three key data engineering challenges.

1. The vectorization challenge

AI systems require enterprise documents to be transformed into vector representations and stored in vector databases for semantic search and retrieval.

Traditional data pipelines were largely designed for structured information such as transactions and operational reports. They are often less suited to unstructured content, including emails, policies, contracts, chat transcripts, and knowledge articles.

As AI deployments grow, organizations need processes that continuously ingest, cleanse, enrich, and vectorize information. Delays can leave AI systems working with outdated data.

2. Data drift and quality degradation

Enterprise systems continually change. New fields are added, formats evolve, and business processes are modified. These changes can introduce inconsistencies that go undetected, leaving AI systems with incomplete or incorrectly mapped data.

Over time, performance can decline. What appears to be an AI problem may actually be a data quality issue originating in the pipeline. Continuous monitoring of data quality, schema changes, and governance is therefore essential.

3. The lack of semantic lineage

In regulated sectors such as banking, insurance, and healthcare, organizations need to understand which information influenced an AI-generated recommendation.

Traditional data lineage tracks where information originated and how it moved between systems. AI environments increasingly require semantic lineage, which shows how data was transformed, interpreted, and used to generate a response.

Without this visibility, compliance, auditability, and risk management become more difficult.



Prem Chand Kumar
CTO, Punjab and Sind Bank

Moving toward a data-centric AI strategy

Organizations often approach AI by focusing on the latest foundation model or comparing LLM providers. A more sustainable approach is to strengthen the data foundation first.

- Invest in platform engineering before chasing AI hype. Modern data platforms, real-time integration, metadata management, and automated quality controls can help AI systems access current, trusted information.
- Treat data as a product. Data requires version control, testing, monitoring, and lifecycle management, much like software.
- Build standardized ingestion pathways. Common ingestion and integration frameworks can help organizations onboard new AI use cases without creating additional silos or bespoke data flows.

The real competitive advantage

Foundation models are becoming increasingly accessible, making trusted enterprise data a potential differentiator. Organizations that can provide models with timely, well-governed information may be better positioned to derive value from AI initiatives.

The AI conversation often begins with models, but sustainable success begins with data. Before asking whether the AI is ready, organizations should ask whether their data foundation is ready. That is where much of the challenge lies. ■

Enterprise AI success depends more on workforce design than model choice



Kyndryl report finds that skills gaps, role redesign, and organizational change are becoming bigger barriers to enterprise AI success than technology itself

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

For the past few years, enterprise conversations around artificial intelligence have centred on models, infrastructure, copilots and more recently agentic AI. The latest Kyndryl 2026 People Readiness Report suggests that while AI adoption has accelerated across enterprises, workforce readiness has moved in the opposite direction.

The findings indicate that AI has already moved beyond experimentation. Nearly 16 oragnisations report broad or embedded AI deployment, while over three-quarters have scaled generative AI across multiple business functions. More importantly, 59% of leaders now view AI as either central to competitive differentiation or a key driver of competitive advantage.

Yet despite this progress, many organisations are struggling to convert AI investments into business outcomes. Only 32% of respondents reported achieving one of their two most desired AI outcomes, while just 11% achieved both.

The readiness gap

One of the significant findings of the report is the widening disconnect between technology readiness and workforce readiness. While 35% of leaders believe their IT infrastructure is ready to support AI, only 23% say the same about their workforce. Organisational culture fares little better at 25%. More concerning is the fact that workforce readiness has declined by six percentage points compared with the previous year.

This concern is reflected in executive sentiment. Nearly 80% of respondents believe AI adoption is advancing faster than their organisations can adapt their workforce, governance structures and operating models. Moreover, leaders expect people-related issues such as workforce skills, role redesign and organisational change to take longer to solve than infrastructure or technology challenges.

AI adoption is advancing faster than organizations can adapt for their workforce, governance structures, and operating models.

Skills as the primary bottleneck

Skills gaps rank among the top obstacles to AI execution, ahead of many technical concerns. More than half of organisations report increasing difficulty in finding external talent with AI-related skills, while 94% believe upskilling existing employees will ultimately be more effective than relying solely on new hires.

Despite recognising the importance of skills development, few organisations have established the mechanisms required to manage workforce transformation at scale. Only a third maintain an accurate inventory of employee skills. Fewer than one-third have proactive upskilling strategies, and only a quarter have created formal career transition pathways for employees whose roles may be affected by AI.

Human-AI collaboration

Contrary to widespread concerns around job displacement, the report suggests organisations are focusing more on redesigning work than eliminating it.

More than half of surveyed organisations have already begun redesigning roles to accommodate AI. Most leaders view AI as a collaborative partner rather than a replacement for human labour. An overwhelming majority believe future roles will require human-AI collaboration, while many organisations are creating entirely new positions dedicated to managing AI workflows, outputs and governance.

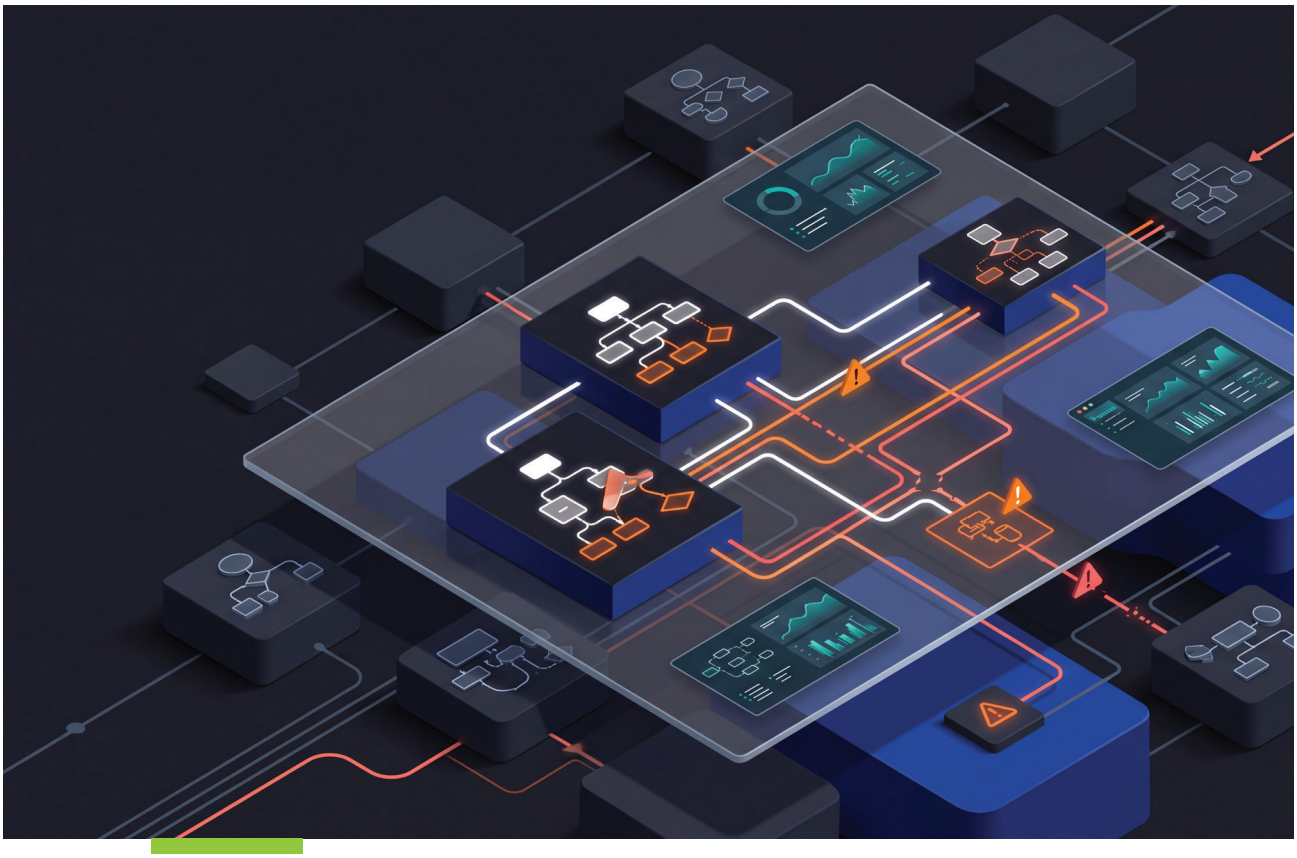
Perhaps, executives still expect humans to remain the primary source of value in areas requiring judgment, ambiguity management, accountability and relationship-building. AI is expected to dominate data analysis and process execution, but decision-making authority remains firmly linked to human oversight

The pacesetter

Among the organisations surveyed, only 9% qualified as “Pacesetters”—companies that have simultaneously redesigned roles, implemented formal change management and achieved workforce readiness.

What distinguishes these organisations is not superior technology. Instead, they have invested heavily in workforce planning, skills visibility, role redesign, change management and governance.

The results are tangible. Pacesetters report significantly stronger outcomes in innovation, revenue growth, risk management and infrastructure modernisation. They also demonstrate substantially higher trust in autonomous AI systems and customer-facing AI deployments. ■



Why observability is the key layer in agentic AI

Agentic AI needs more than autonomy to scale. Observability builds trust, strengthens security and makes AI agents accountable

By **Bianca Lewis** | editor.tech@timesgroup.com

As AI agents move from experiment to enterprise, companies are progressing from chatbots handling FAQs and models summarizing documents to agentic AI systems that can now reason across multiple steps, select and use tools, access live data, and execute actions with minimal human intervention. But how can you verify what the agent is doing, ensure it's done correctly, and intervene when it isn't?

Unlike traditional software, where a failure can be traced to a specific line of code or infrastructure fault, agentic AI operates through chains of reasoning and decision-making that aren't predetermined and aren't always reproducible. An agent that can act on your behalf across systems, data sources, and workflows is only as reliable as your ability to observe what the system is doing. Trust has to be built, not assumed. In the age of agentic AI, observability is the trust layer that determines whether autonomous systems can be deployed

at scale, governed responsibly, and relied upon when it matters.

When the reasoning is the process

To understand why observability matters so much here, it helps to understand what makes agentic AI different from everything that came before it. Traditional software systems follow deterministic paths. Early AI systems like recommendation engines and large language models used as assistants were auditable in a similar manner: a model received an input, produced an output, and the relationship between the two could be inspected.

Agentic AI breaks this pattern. An agent given an objective may use external tools, query multiple data sources, spawn sub-agents, and produce outputs through a chain of reasoning that is neither predetermined nor always reproducible. The reasoning itself is the process. And that process has historically been invisible. For enterprises deploying agentic AI, that invisibility is a problem observability solves.

Observability expands its mandate

For years, observability focused on monitoring CPU utilization, memory consumption, and application latency. And while that foundation remains necessary, for agentic AI, it is no longer sufficient. Organizations deploying AI agents need visibility into an entirely new set of questions: what data sources did the agent consult? Where did the agents get their context from? What tools did it invoke, and in what sequence? What reasoning path led to that output? Did it act within the boundaries it was given? These questions are the ongoing evidence that makes agentic AI governable and trustworthy enough to deploy in production. This is where a platform like OpenSearch becomes critical. OpenSearch acts as the secure memory layer for these agents to trace historical context. By capturing trace data, tool execution logs,



Bianca Lewis
Executive Director, OpenSearch Software Foundation

and system telemetry (usually via OpenTelemetary) into a centralized, highly searchable datastore, organizations can instantly query the entire lineage of an agent's decision-making process—turning invisible AI reasoning into predictable, auditable logs.

Building the evidence trail

Trust emerges from transparency and observability is how that transparency gets built. When stakeholders can review and comprehend how an AI system behaves at scale. Without that evidence trail, governance is aspirational.

This dynamic is familiar. Organizations were initially reluctant to move critical workloads to the cloud out of fear of losing visibility and control. Advanced observability tools bridged that trust gap by providing deep insight into distributed environments.

Agentic AI is following a similar arc, and the organizations that instrument their AI systems now will be the ones that can move fastest ahead of future regulatory changes. Auditability, explainability, and accountability are becoming mandatory business requirements rather than optional capabilities. To build this evidence trail, enterprises are increasingly leaning on OpenSearch as the core architecture for AI governance. Because it natively indexes both structured system metrics and unstructured text traces at scale, it provides a tamper-evident ledger of every API call, prompt variation, and sub-agent action. When compliance or business logic demands an explanation for an agent's output, OpenSearch allows teams to reverse-engineer the entire chain of execution in seconds.

Security in an age of autonomous agents

The same visibility that makes agentic AI governable also makes it defensible. AI agents operate across a broad and expanding footprint of tools, APIs, databases, and services that may themselves have access to sensitive systems. Without observability, security teams are effectively blind to what agents are doing in real time. Abnormal behaviour, unauthorized access attempts, excessive privilege usage, and potentially damaging action sequences can escalate into serious incidents before anyone notices. With comprehensive agent telemetry, security teams can detect, investigate, and respond to threats with the same rigor they apply to traditional infrastructure. As AI agents are given more autonomy, ongoing visibility into their actions becomes non-negotiable.

The trust layer

Every enterprise deploying agentic AI will eventually face the same question: not whether their agents are capable, but whether they can be trusted. Capability is increasingly available to everyone. Trust is not. It has to be built deliberately, through systems and practices that make autonomous behaviour visible, auditable, and governable.

Organizations that treat observability as a strategic foundation now, not just a diagnostic tool to reach for when something goes wrong, will be the ones that can deploy agentic AI with confidence, scale it responsibly, and defend it when it matters. The future of AI is agentic. The future of trust is observable. ■

Why trustworthy AI needs a measurement standard

As AI reshapes industries at breakneck speed, **CSIR-NPL Director Prof. Venu Gopal Achanta** makes the case that metrology — the 150-year-old science of measurement — holds the key to certifying, benchmarking, and building trust in AI systems

By **Jagrati Rakheja and Jatinder Singh** | editor.tech@timesgroup.com

As artificial intelligence reshapes every industry at breakneck speed, one question often goes unanswered: who verifies that AI works the way it claims? Prof. Venu Gopal Achanta, Director of CSIR-National Physical Laboratory (CSIR-NPL), believes metrology, the science of measurement, holds the answer.

In a recent conversation with ET Edge CIO&Leader, he draws parallels between India's 150-year-old metrology tradition and the emerging need for AI certification, benchmarking, and traceability. He unpacks why quantum sensors, not quantum computing, deserve India's immediate attention, what the German and American models offer as blueprints, and why semiconductors should top India's technology priority list over the next decade.

Achanta argues that India's path to becoming a technology powerhouse by 2047 runs through its most invisible science: precision measurement. Backing that vision is a sizeable institution: CSIR-NPL runs on a team of around 800 people — about 145 scientists, more than 150 technical staff, and about 100 people in administration, finance, and support roles, plus many contractual staff and students, including roughly 200 industry-sponsored PhD scholars.

Excerpts from the interaction:.

CIO&Leader: Today, NPL is the custodian of India's measurement standards. In an AI-first world, could NPL play a similar role for trustworthy AI through standards, benchmarking, or certifications? Could India eventually require an AI equivalent of calibration for critical applications?

PROF. VENU GOPAL ACHANTA: This is a very important question, and it is pertinent for the future. In metrology, we have the seven SI units, and these, along with derived units, underpin every measurement. Scientific metrology, which measures these seven SI units as precisely as possible, remains important even after 150 years, and going forward, AI will both influence and require metrology.

AI's implications depend on where it's applied. In biomedical applications, for instance, ethical issues arise, how do we distinguish between generative AI, data analysis, and other approaches? These are things we need to address.

In the metrology community, we're working on digital transformation. At the CIPM level, there's a body called Forum MD, where countries work together to digitise how we handle data, digitise measurements, and bring AI into data analysis and other applications. We are using AI while also building metrology models to ensure the data source is reliable, for AI, one of the most critical things is ensuring the integrity of the source.

So we are working on a network of sensors and the metrology behind them. Calibration, testing, and certification of various kinds of sensors is something we're working on together as a metrology community.

CIO&Leader: AI is evolving so fast that it's reshaping how everything works, making responsible AI an urgent priority for CTOs and innovation leaders at large



"The clearer focus should be on quantum sensors, not on the buzz around quantum computing. Let the quantum computers come when they come."

enterprises. Where do we need to invest more now to ensure students and researchers are equipped to build trustworthy AI models, and how do you see the market evolving for responsible AI from here?

PROF. VENU GOPAL ACHANTA: I'll point to the German model. The German government mandated PTB Germany, their national metrology institute, equivalent to NPL India — to develop digital transformation models, starting with certification, that would become a standard across all industries. They take use cases from each industry and implement a common base platform on which industry-specific verticals are built. That mandate came from the government side.

Unfortunately, metrology is still not very pervasive here. Quality is something the Prime Minister keeps emphasising — our products should be seen as high-quality, high-value products globally. Metrology is the backbone of that; global acceptance of products comes from the certification metrology institutes provide.

The MRA convention dates back to 1875, and under the CIPM Mutual Recognition Agreement, 92 countries have signed on. A certificate issued by NPL India is accepted in 92 countries worldwide.

We need to strengthen our metrology activities by bringing in industry, R&D institutions, and universities. NIST in the US is a good example — they fund people from universities to come work in their labs, develop new techniques and methods, and sometimes translate them into technologies. We need that kind of platform. We still need dedicated funding for metrology.

CIO&Leader: Apart from funding, do you think India has the talent required to become a key powerhouse in setting these standards?

PROF. VENU GOPAL ACHANTA: We have the workforce in abundance, but government labs need to strengthen their staffing — many posts have been reduced over the years. These are highly skilled jobs; people train and upgrade their skills over decades to become master metrologists. India has talent, but its government labs need stronger staffing — not just scientists, but technical officers and technical assistants too. We have a talent pool within the country. We need to use it.

CIO&Leader: What are some of the invisible scientific capabilities — whether in measurement science, quantum, or advanced instrumentation — that India should invest in to lead the next wave of innovation?

PROF. VENU GOPAL ACHANTA: Quantum is one area where we're already doing well. It spans quantum computing, quantum communication, quantum sensors in metrology, and quantum materials — and traditionally, we're very strong in materials.

Quantum computing is different, and only a few groups are working on it, whether superconducting or neutral-atom-based, under the National Quantum Mission. But the low-hanging fruit is quantum sensors, where sensitivity can be much higher and precision can reach very high levels for different applications. The clearer focus should be on quantum sensors, not on the buzz around quantum computing. Let the quantum computers come when they come — before that, AI, high-performance computing, and the modest quantum computing resources we have today should be pooled together, and we should develop utilities around them.

For quantum sensors, we have strong facilities at CSIR-NPL. Over the last four to five years, we've invested in them — we developed and demonstrated one of India's first magnetometers here and transferred it to a startup for strategic applications. We're now transferring technologies to industry across the quantum sensors vertical: magnetic field, electric field, RF field sensing, single-photon sources, and quantum random number generators. This includes calibration and certification facilities for all of these.

CIO&Leader: Which countries are we transferring our knowledge to, and with whom are

we closely aligning as a global powerhouse?

PROF. VENU GOPAL ACHANTA: This is tricky, because many countries are closing their doors, they don't want to transfer knowledge or share cutting-edge technologies. In metrology, there's a body called the National Metrology Institutes of Quantum (NMIQ), which includes the US, Japan, Australia, Canada, Germany, and the UK. These countries want to develop standards for these futuristic technologies among themselves, and that's where we're working to be part of the conversation.

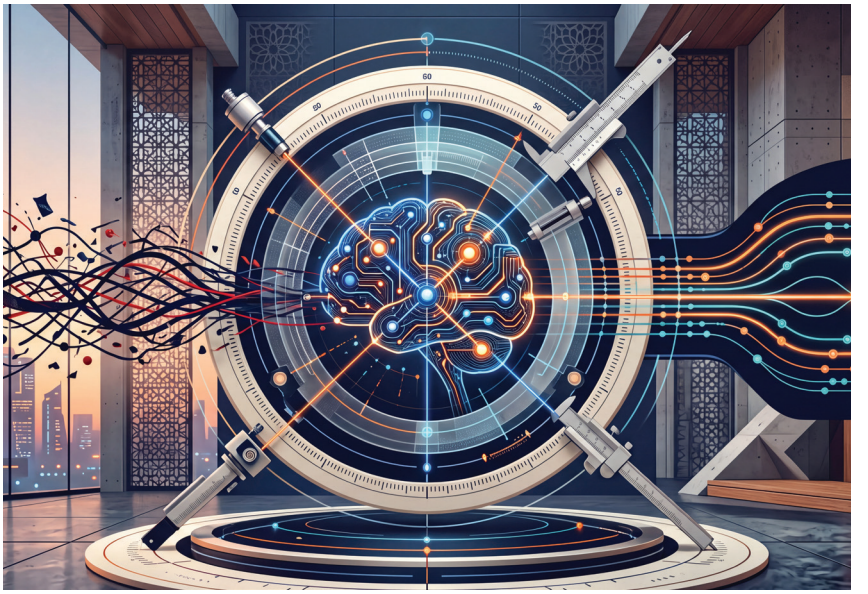
We now have about 150 quantum startups in the country, including 40-plus in the Amaravati Quantum Valley, and we're working with them. The idea is to set up the primary national standard here at NPL India and compare our measurement facilities with those worldwide, so that a certificate issued by NPL is accepted globally. We also want to extend this to Amaravati and other places, so certification and standards facilities are available for startups there too.

We still collaborate closely with PTB Germany — recently on energy and renewables, whether hydrogen or solar. Their president visited recently, and we signed an MOU focused on energy materials, the energy sector, metallurgy, and fossil energy.

CIO&Leader: Which emerging technology or scientific discipline do you believe India is significantly under-investing in today, despite its potential to shape our future?

PROF. VENU GOPAL ACHANTA: Quantum has considerable visible funding, but translation isn't happening — only a few groups get picked up and funded. Many industries get funds based on idea alone. For the past three years, I've been trying to set up standards here, and we haven't received dedicated funding for it.

People need to focus on metrology. Scientific leaders should support it, because it will eventually be essential for every product coming out of India. Unless we provide metrology support, startups and industry will have to go abroad to certify or characterise their devices.



"R&D institutions so far have focused on publications; now we should focus on translating that knowledge and know-how into industry... so innovation becomes influence."

We do have an AI mission, a quantum mission, a semiconductor mission, and a high-performance computing mission — this government is willing to invest heavily in each of these. But we also need to focus on the supporting metrology behind them.

CIO&Leader: Government research institutions are often seen as centres of scientific excellence. How can institutions like NPL and CSIR collaborate more effectively with start-ups and industry to accelerate India's deep-tech ecosystem?

PROF. VENU GOPAL ACHANTA: I am part of the Quantum Valley initiative, and we are making use of that. We are organising a hackathon there, because we need more interaction. It's a big country, unlike a European country like Germany or Italy, one of our states is like a continent in scale, so whatever we do, however much publicity we give it, there's always more need.

Minister Dr Jitendra Singh proposed the idea of "One Week, One Lab," we bring in stakeholders, students, and industry, one group each day. The resources needed for this are huge.

There is also an ongoing sense of competition globally, the US has dominated so far, but the Chinese presence in this space has grown quickly, and it's still evolving; we don't yet know which approach will be most widely accepted. Industry's caution may also stem from the continued need for manual verification. Generative AI, NLP, and similar technologies are advancing quickly, but at this stage I still take it with a pinch of salt — I would still want to check whether what it's producing is right.

CIO&Leader: If India had to place one major technology bet over the next decade, AI, quantum,

semiconductors, or scientific infrastructure, which would you prioritise, and why?

PROF. VENU GOPAL ACHANTA: Semiconductors, first and foremost, because semiconductor technologies drive many of these other fields, the fabrication facilities built for semiconductor technology are already helping the quantum world today. Let's invest in semiconductor technologies first and then translate that investment to other areas.

Quantum should be the second priority, we need to be self-sufficient there and not simply buy these capabilities from others. Semiconductor technology requires significant investment in scientific infrastructure, which is largely in place in the country already but needs expansion. From a government perspective, we have a high GDP, but the population we're dividing it across is large, so priorities matter.

I would prioritise semiconductors and related technologies, then translate that into quantum. AI will follow, because to me AI is still an elite resource, but it will eventually span every vertical.

CIO&Leader: As we look toward Viksit Bharat 2047, what would define success for India as a technology nation, not just in innovation, but in scientific leadership and influence?

PROF. VENU GOPAL ACHANTA: Coming from a metrology institute, I have to say: metrology. In any top economy, metrology is directly linked to economic strength, the two go hand in hand. Quality depends on how well we characterise and certify our own equipment, and metrology underpins whatever comes next, including defence.

The key shift is translation. R&D institutions so far have focused on publications; now we should focus on translating that knowledge and know-how into industry, that's the government's push too. We should work toward that until 2047, so innovation becomes influence. We should aim to cut our imports of high-tech equipment by 70–80 per cent, and make high-end scientific equipment ourselves, in India.

CIO&Leader: In terms of tech-focused priorities for your organisation over the next two to three years, could you share some of the top areas of focus?

PROF. VENU GOPAL ACHANTA: The seven base SI units remain the starting point for everything we do.

Over the last five years, we've expanded our scope to include traditional medicine. How do we bring Ayush and Siddha medicine under the SI system of metrology? Some of these traditional measures are still imprecise, a "pinch" of an ingredient, for instance, so we are working with the Ministry of AYUSH, Siddha, and our sister labs to develop phytochemical standards and certified reference materials within the country. These will be the standards against which the purity of these medicines is checked, helping harmonise the entire traditional medicine ecosystem.

On the renewable energy side: about four years ago, we asked what challenges we as metrologists should address. One of the biggest is grid stability. There are more than 10,000 phasor measurement units across the grid, measuring current, voltage, and phase over time to monitor how a fluctuation at one location affects another. Precise time-stamping is essential for this, so we needed to ensure all these units were calibrated and reporting consistent numbers. We now have a robust system where every PMU in the country is traceable to the NPL facility.

The government is also targeting 50 gigawatts of solar installation per year. A 1 per cent error in estimating a solar panel's efficiency translates to roughly a billion-dollar difference in investment, so we need to measure the quantum efficiency of solar cells precisely, we've set up a facility for that.

Looking 15–20 years ahead, these solar panels will become waste. Recycling efforts so far have focused on extracting precious metals, silver, gold, platinum plus aluminium frames, glass, and epoxy. We asked what happens to the silicon itself, and we've developed a process to recycle and reuse it.

All of this reflects the government's priorities and our effort to move from innovation to impact in nation-building. Quantum is another area where we're contributing in the same way.

CIO&Leader: Are there any internal AI tools you are developing to enhance your research capabilities and ensure researchers across the ecosystem benefit from the knowledge generated?

PROF. VENU GOPAL ACHANTA: We do mostly AI and machine learning models. One recent example, done with a student, involved Jantar Mantar. These yantras were built in the 1700s, so they're over 200 years old, and we wanted to understand their performance. The student spent five days manually measuring and recording data, then comparing it against NTP (atomic clock time). We found that once you correct for Earth's rotation and other effects, the instruments perform as expected, with a best resolution of two seconds.

For that, we needed to build models. Indian Standard Time is based on a bank of atomic clocks, each with its own uncertainty and fluctuation that varies through the day, so we take an ensemble average across all of them, and we are developing AI models to improve the accuracy and sensitivity of that averaging.

We should be proud that of the roughly 400 optical clocks worldwide contributing to global time, what used to be called GMT and is now UTC, India's clocks have the sixth-highest weightage in the world. It's one of the best contributions we make. ■

Automate the service standard, preserve human judgment

Anand Mahalingam, Vice President – Data Science, Digit Insurance, on how bounded autonomy, real-time decisioning and agentic AI are reshaping claims, underwriting and customer experience

By **Jatinder Singh** | jatinder.singh1@timesgroup.com

Insurance is entering a new phase of AI adoption. The focus is shifting from using AI to accelerate individual processes to building intelligent systems that can interpret data, make decisions and act in real time. From claims and underwriting to fraud detection and customer engagement, agentic AI is beginning to reshape how insurers operate and how customers experience insurance.

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.

In this conversation with ET Edge CIO&Leader, Anand Mahalingam, Vice President – Data Science, Digit Insurance, discusses the organisation's approach to bounded autonomy, real-time claims and underwriting, AI governance, and the emerging use cases that could redefine the insurance experience.

CIO&Leader: How is Digit approaching the shift from AI-assisted workflows to autonomous decision-making systems?
ANAND MAHALINGAM: At Digit, our approach to autonomy is evolutionary rather than abrupt. We are moving from AI-assisted workflows towards bounded autonomous systems, where AI operates within clearly defined decision frameworks, while humans retain oversight over high-impact decisions.

AI systems are already making autonomous decisions in areas such as document routing, processing pathways and risk identification.

In claims processing, for instance, multiple steps—including document validation, risk screening and duplication checks—can be executed automatically and in parallel. This significantly reduces manual intervention while improving consistency and turnaround times.

We are also extending agentic principles beyond core operations. AI-powered conversational and calling agents are increasingly managing customer interactions across lifecycle touchpoints, enabling personalised engagement at scale.

Looking ahead, we see AI playing a larger role in straight-through claim approvals, autonomous fraud triage and intelligent escalation. However, in insurance, autonomy must always be paired with accountability, transparency and human oversight. The objective should be to automate the service standard while preserving human judgment where it matters most..

CIO&Leader: What role is AI playing in enabling scalable, real-time claims and underwriting processes?
ANAND MAHALINGAM: AI is playing a pivotal role in how we scale operations at Digit, particularly across claims and underwriting. We have been moving from sequential processing towards real-time, event-driven decisioning.

In claims, AI enables document classification, information extraction, validation and fraud checks to run simultaneously rather than sequentially. This allows us to handle large volumes more efficiently while significantly reducing turnaround times.

High-confidence cases can be processed automatically, while complex or ambiguous cases are routed for human intervention. This creates a balance between speed and accuracy.

On the underwriting side, AI-driven models enable us to assess risk at a much more granular level and in near real time. Processes that traditionally took hours can increasingly be completed in minutes as structured and unstructured signals are evaluated simultaneously.

This capability becomes particularly important as insurance is increasingly embedded into real-time customer journeys through APIs and partner platforms. The ability to make fast, contextual decisions is becoming a key differentiator.



"AI is playing a pivotal role in how we scale operations at Digit, particularly across claims and underwriting. We have been moving from sequential processing towards real-time, event-driven decisioning."

CIO&Leader: As AI systems become more agentic, what are the key governance, risk and security considerations?

ANAND MAHALINGAM: As AI becomes more agentic, governance cannot be an afterthought; it has to be embedded into the architecture from the outset. At Digit, our focus is on ensuring that every AI-driven decision is transparent, traceable and auditable.

Each decision should have a clear trail of inputs, outputs and confidence levels. This allows us to review outcomes internally while also supporting regulatory and compliance requirements.

Equally important is maintaining a strong human-in-the-loop framework. Automation can handle standard and high-confidence scenarios, but critical decisions—particularly around complex claim approvals—require human judgment where necessary.

From a risk and security perspective, we operate within defined data-governance boundaries. Sensitive customer information needs to be anonymised, securely processed and used responsibly. As we move towards more agent-driven architectures, we are also placing greater emphasis on system-level controls that define roles and permissions across AI agents, monitor feedback loops and ensure that autonomy remains bounded.

Ultimately, increased autonomy should lead to greater transparency and trust—not opacity.

CIO&Leader: What are the emerging use cases where AI could drive the next wave of innovation in insurance?

ANAND MAHALINGAM: The next wave of insurance innovation will be driven by AI's ability to enable real-time, embedded and event-driven insurance experiences.

Take travel insurance, for example. AI can enable faster and more frictionless payouts when customers face disruptions such as flight delays or cancellations. As data ecosystems mature, we could see more use cases where payouts happen almost instantly, without conventional claims-processing journeys.

Embedded insurance represents another significant shift. Insurance can become part of the customer journey rather than a standalone purchase. Across travel, mobility and financial services platforms, AI can enable real-time underwriting, contextual pricing and instant policy issuance, allowing insurance to become seamlessly integrated into everyday transactions.

Conversational AI is also evolving from a support tool into a decision interface. Instead of simply answering customer queries, AI agents could guide customers, explain policies and assist with decision-making.

Autonomous claims settlement is another area with significant potential. AI systems could eventually validate, assess and approve standard claims with minimal human intervention. Combined with advances in real-time underwriting and continuous risk monitoring, this could move insurance towards a model that is more proactive, personalised and responsive.

The larger opportunity is to move from insurance that responds to events to insurance that can anticipate and act on them. ■

The network is now the foundation for business transformation

Sajan Paul, General Manager, HPE Networking India explains how AI-native, self-driving networks, built-in security, liquid cooling, AI grids and autonomous operations will help enterprises simplify complexity while preparing infrastructure for the AI era

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

As artificial intelligence reshapes enterprise IT, networking is undergoing one of its most significant transformations since the advent of cloud computing. AI workloads, hybrid environments, distributed applications and an expanding threat landscape are pushing enterprise networks beyond the limits of manual operations. At the same time, Indian organisations are grappling with infrastructure complexity, cybersecurity risks, talent shortages and sustainability challenges as they modernise for the AI era.

In this interview Sajan Paul, General Manager, HPE Networking India, discusses HPE's vision for autonomous networking, why security must be embedded into network architecture from the ground up, how the integration of Juniper and Aruba is shaping the company's networking strategy, and what enterprise networks could look like over the next five years.

CIO&Leader: HPE has spoken extensively about self-driving networks. What does this vision mean in practice for enterprise CIOs?
SAJAN PAUL: Networks have evolved far beyond being simple connectivity layers. Today, they are the strategic foundation for digital transformation, hybrid cloud, Industry 4.0 and AI-driven business models.

That transformation has also introduced enormous complexity. There are now more than 35 billion connected devices worldwide, and every new branch, IoT deployment, cloud application or remote user increases operational complexity. Enterprises depend on networks for mission-critical operations while simultaneously facing networking skill shortages and an expanding attack surface.

Managing this complexity manually is becoming unrealistic.

That's why HPE is building AI-native networking, where AI becomes a foundational capability rather than an add-on. AI continuously analyses network behaviour, predicts failures, detects anomalies and optimises performance in real time.

The next step is the self-driving network.

Our vision is to build networks that continuously learn, optimise, heal and increasingly operate autonomously across the enterprise, not just for isolated use cases but across the complete networking environment. This remains one of HPE Networking's highest strategic priorities globally.

CIO&Leader: As AI expands the enterprise attack surface, how should organisations rethink network security?
SAJAN PAUL: I prefer calling it secure-native networking.

Just as organisations have embraced cloud-native and AI-native architectures, security must also become native to the network. It can no longer



"Security has to be thought through from the ground up, right from designing hardware to designing a software stack, and then through its entire life cycle."

be treated as a layer added after deployment because threats evolve far too quickly.

Security has to be designed into every layer, from hardware architecture and software development to the operational lifecycle.

Take a network with one thousand devices. If a critical vulnerability is discovered, organisations cannot afford months to deploy patches. That's why HPE has adopted a microservices architecture.

Instead of relying on monolithic software, every networking function operates independently. If one component is affected, it doesn't compromise the entire system, and individual services can be updated rapidly across thousands of devices without replacing the entire software stack.

The second pillar is Zero Trust.

We assume no network is inherently trusted. Internally, we describe this as "Coffee Day networking"—whether you're in your office, working remotely or connected through public Wi-Fi, the security posture should remain identical.

This philosophy underpins our Secure Edge portfolio, SD-WAN, Secure Service Edge (SSE) and Zero Trust Network Access (ZTNA) architecture.

We're also preparing for the post-quantum era. Quantum computing could fundamentally weaken today's encryption methods, so HPE Labs is already investing in quantum-safe security architectures to help organisations prepare before the threat becomes reality.

CIO&Leader: HPE recently announced deeper integration between Juniper Mist and Aruba Central. What does this mean for customers?

SAJAN PAUL: Our guiding principle throughout the Juniper acquisition has been simple: no customer should be left behind.

Both Aruba and Juniper have strong customer communities, and our objective is to preserve those investments while creating additional value through integration rather than forcing migrations.

An important advantage is that both Aruba Central and Juniper Mist were built on cloud-native, microservices architectures. That common foundation allows us to introduce new AI capabilities rapidly across both platforms.

Rather than replacing one platform with another, we're making the customer experience increasingly consistent.

A useful analogy is Android and iOS. Different operating systems, but users expect similar application experiences. That's exactly the direction we're taking.

Capabilities available in Mist, particularly AI-driven operations, will progressively become available in Aruba Central. Likewise, Aruba's strengths in identity services and customer intelligence will become available in Mist.

Hardware flexibility is equally important.

Customers often ask whether they should choose Aruba or Juniper hardware. Our answer is that whichever platform they choose, it will continue to be fully supported.

We've even introduced a dual-personality access point capable of operating with either Aruba Central or Mist.

Ultimately, by combining the strengths of both portfolios, we believe we're creating one of the industry's most comprehensive platforms for AI-driven, self-driving network operations.

CIO&Leader: AI workloads have introduced entirely new networking requirements. How is HPE supporting hyperscalers and private cloud operators?

SAJAN PAUL: There are really two conversations today.

The first is AI for Networking, where AI automates and simplifies network operations.

The second is Networking for AI, which is about building infrastructure capable of supporting increasingly demanding AI workloads.

Modern AI generates enormous east-west traffic between GPUs, storage and compute clusters. Networking performance directly

affects AI training times and inference speed.

The key metric is no longer simply bandwidth—it's job completion time.

At HPE Discover, we introduced our Broadcom Tomahawk 6-based platform delivering 1.6 terabits per second per port and 102.4 terabits per second of switching capacity

Equally important, these switches are fully liquid cooled.

Liquid cooling reduces power consumption, increases rack density and enables data centres to maximise increasingly scarce space. Today, operators even evaluate AI infrastructure using metrics such as tokens generated per square foot.

Beyond hardware, software orchestration is equally important.

We've integrated our Apstra Intent-Based Networking platform to automate network deployment and lifecycle management while maintaining operational consistency across large AI environments.

We're also expanding operating system flexibility by supporting both Junos OS and SONiC, the open-source networking operating system increasingly adopted by hyperscalers.

Whether organisations deploy NVIDIA clusters, AMD infrastructure or private AI environments, the network should never become the bottleneck. It should accelerate AI training and inference while simplifying operations.

CIO&Leader: India continues to face challenges around power, sustainability and networking skills. How can enterprises address these constraints while scaling AI infrastructure?

SAJAN PAUL: Scaling AI infrastructure isn't simply about deploying more GPUs. Organisations must optimise power, cooling, utilisation and connectivity together.

Liquid cooling is one of the most significant technologies enabling this transition. Drawing from HPE's supercomputing expertise, we've adapted liquid cooling into networking to improve thermal efficiency, reduce power consumption and minimise infrastructure footprint.

"Security has to be thought through from the ground up, right from designing hardware to designing a software stack, and then through its entire life cycle."

Sustainability also depends on utilisation.

In reality, network switches rarely operate at full capacity continuously. Traffic fluctuates, leaving processing resources idle.

We've therefore developed AI-driven sustainability algorithms that monitor hardware utilisation and intelligently power down unused processing components without affecting performance.

Beyond individual devices, enterprises must rethink infrastructure deployment itself.

Finding abundant power, cooling and connectivity in a single location is becoming increasingly difficult. Instead of building one massive AI factory, organisations are creating multiple AI facilities distributed across different locations and connecting them through ultra-low-latency networking.

We refer to this as the AI Grid. Several AI factories, connected through high-speed networking, can function as a single intelligent infrastructure while providing greater resilience and scalability.

HPE is investing heavily in these ultra-low-latency interconnects and advanced traffic engineering to make distributed AI infrastructure practical.

CIO&Leader: Looking ahead, what will enterprise networking look like five years from now?

SAJAN PAUL: TWe're much closer to autonomous networking than many people realise.

A few years ago, self-driving cars seemed futuristic. Today, they're already operating commercially in several parts of the world.

I believe self-driving networks will follow a similar trajectory. Given the pace of innovation, many of these capabilities could become mainstream within the next three years.

The self-driving network will become an imminent reality. I'd be surprised if it will happen in less than three years.

Of course, autonomous networking depends on continuous learning.

AI systems are only as good as the operational data used to train them, which is why HPE continues investing heavily in improving model accuracy, validating outcomes and enhancing AI explainability.

Transparency will become increasingly important as enterprises need confidence in how AI systems make decisions.

We're already seeing humans shift from operational execution towards architecture and governance.

Just as software developers increasingly act as system architects rather than writing every line of code, network engineers will spend less time performing repetitive operational tasks and more time designing policies, architectures and governance frameworks.

That said, I don't believe humans will disappear from the process. Compliance, governance and critical decision-making will continue to require a human-in-the-loop.

Looking further ahead, the boundaries between servers, compute, networking, storage and software will continue to blur.

Self-driving capabilities won't remain confined to networking—they will increasingly extend across the entire enterprise technology stack, creating infrastructure that is autonomous, intelligent and capable of managing itself while allowing IT teams to focus on higher-value strategic outcomes. ■

AI is becoming indispensable to cyber defence

Fabio Fratucello, Vice President and Field CTO Worldwide, CrowdStrike explains why identity is becoming the new cyber battleground, urging enterprises to secure AI agents with just-in-time access, unified visibility and AI-driven detection and response

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

As organisations increasingly adopt cloud, AI and autonomous systems, identity has emerged as a critical battleground in cybersecurity. Attackers are increasingly using valid credentials to move through enterprise environments, while the rise of AI agents is creating a new class of non-human identities with access to sensitive resources. At the same time, shrinking attack timelines are making human-led detection and response increasingly difficult.

In this interview, Fabio Fratucello, Vice President and Field CTO Worldwide, CrowdStrike, discusses why identity has become an attractive entry point for attackers, how enterprises can secure AI agents through just-in-time and just-enough access, why accountability for autonomous systems still rests with humans, and how organisations can build a unified, cross-domain security architecture capable of responding at machine speed.

CIO&Leader: Adversaries are increasingly choosing to log in rather than break into systems, using genuine-looking employee credentials or non-human identities. What gaps in traditional perimeter and endpoint security are making identity an easier entry point for attackers?

FABIO FRATUCELLO: Traditionally, adversaries targeted endpoints because that is where runtime exists and where there is user interaction. A person is typing on the keyboard, making the endpoint an ideal environment to target.

However, EDR technology has made the endpoint more difficult to compromise. As a result, adversaries have started looking for other opportunities where they can gain an advantage.

That is where identity comes in. Identity-driven attacks, malware-free attacks and compromised identities have been trends we have seen for several years. Adversaries are looking to obtain usernames and passwords because they can then log in rather than break in.

From a defensive perspective, logging in with a valid credential generates less noise. There are no brute-force attempts or authentication failures to immediately trigger defensive mechanisms. While behavioural signals may still raise concerns, detecting and responding to an attacker using a legitimate identity becomes much more challenging.

CIO&Leader: When attackers blend into normal operational traffic using valid credentials, how can security teams distinguish between legitimate activity and a compromised identity in real time?

FABIO FRATUCELLO: Identity threat detection and identity threat protection capabilities can make a significant difference. Authentication is an important part of the puzzle, but it is only one signal.

Logging into a laptop, for example, may simply be the first step in an attack. The adversary may then attempt to move laterally to another laptop, server or system, either on-premises or in the

cloud. Organisations therefore need visibility into the entire sequence and the ability to intervene at that level.

The challenge is that breakout time is becoming shorter. Over the last five years, adversaries have continued to reduce the time required to move through an environment. The last Global Threat Report recorded an average breakout time of 29 minutes, while the fastest e-crime breakout time was less than 30 seconds.

For defenders, that means human-initiated response is increasingly unlikely to be fast enough.

Identity threat detection and protection therefore need to be part of a broader platform capable of understanding the entire attack. We call this a cross-domain attack, where an adversary moves from an endpoint to an identity and then to another logical or physical part of the organisation.

With such short response windows, AI capabilities are also becoming essential because defenders increasingly need machine response time to counter sophisticated adversaries.

CIO&Leader: As organisations deploy autonomous AI agents and non-human identities, what unique security risks do AI agents introduce compared with traditional service accounts or human employees?

FABIO FRATUCELLO: This is one of the major shifts we are seeing. Organisations are embracing AI to deliver faster and better outcomes, and now we have agents that can effectively operate as superhuman employees.

Agents can initiate work and activity, access resources, and use identities to authenticate into different parts of an organisation. The major difference is that they operate 24/7, do not get tired and can process enormous amounts of information in a very short time.

From a defensive standpoint, this makes identity controls even more important.



"Identity threat protection and identity threat detection are key capabilities from a detection and response perspective, but then again, this is one piece of the puzzle."

The fundamental capabilities applicable to human identities also apply to AI identities, but organisations need to go further. Concepts such as just-in-time and just-enough access become critical.

An agent should only be able to access particular resources for a defined period and only have the level of access necessary to perform its task. These controls create guardrails around what an agentic identity can do.

That becomes particularly important as we see emerging attacks such as direct and indirect prompt injection, model tampering and skill abuse. Strong identity controls can limit the potential impact of such attacks.

CIO&Leader: If AI agents can make decisions, access databases and delegate tasks independently, who ultimately owns accountability when an agent breaches a security boundary?

FABIO FRATUCELLO: I don't think there is a straight answer at this point because organisations are likely to handle accountability differently.

However, in the majority of organisations we work with, the person, team or human remains accountable to the C-level and the board. At this point, accountability is not something we have seen organisations being able to pass down to the machine.

That is why governance, visibility and guardrails are becoming central to AI adoption. Leaders need confidence that autonomous systems are operating within defined controls and within the organisation's risk appetite.

The journey generally starts with visibility. Organisations first need to understand how AI is being used internally, including generative AI accessed through browsers and agents operating on endpoints.

Once visibility is established, organisations can introduce governance and guardrails, such as authorising specific services and restricting others.

More mature organisations are then moving towards AI Detection and Response, or AIDR. At that

stage, organisations typically have an AI strategy, dedicated teams and a clearer understanding of their AI adoption. They can then introduce technical controls to protect agents from adversaries, human mistakes, data exfiltration and data leakage.

Ultimately, accountability continues to sit with the security leader, CIO or whoever represents the AI function at the leadership level.

CIO&Leader: Why have static access policies and standing privileges become a liability in the age of autonomous AI? What does continuous, risk-aware authorisation look like?

FABIO FRATUCELLO: Organisations at the forefront have already moved from static privilege to just-in-time privilege, even for human users.

The principle is straightforward. A user may need a particular privilege to perform a specific activity, but once that activity is completed, there is no reason to maintain that level of access. Reducing privileges reduces the potential blast radius if the account is compromised.

The same principle becomes even more important with AI agents because they are effectively superhuman. They can perform activities at much greater volume and speed than a person.

Imagine an agent with domain credentials. If an adversary compromises that agent, they could potentially tamper with it, modify how it behaves and use those credentials to operate at the domain level.

That can become a very serious problem very quickly.

The ideal scenario is for the agent to have access only to limited resources and only the ability to read or modify what is necessary. If the agent is compromised, the attacker is similarly constrained.

That gives defenders more time to detect, respond and contain the incident. The more difficult and time-consuming we make the attack for the adversary, the greater the likelihood that they will make a mistake, generate noise and give defenders an opportunity to respond.

"We need AI that is designed to deliver defensive outcomes, and we need that to be part of the platform to be effective."

CIO&Leader: How can organisations bring human, non-human and AI identities under a unified security architecture without slowing down business operations?

FABIO FRATUCELLO: The problem is becoming increasingly relevant because organisations are using more technologies, generating more data and creating more potential attack domains.

Looking at these domains individually is no longer sufficient. A security tool may provide a good picture of what is happening in one area, but that does not necessarily explain the entire attack or what is likely to happen next.

The best approach is a modern platform that provides cross-domain visibility and AI capabilities for defenders.

The objective should be a harmonised view rather than simply bringing different technologies into one console. Identity telemetry and endpoint telemetry, for example, need to speak the same language and be presented cohesively so defenders can understand the entire attack story.

It is like reading an entire book rather than reading only Chapter 3. If someone logs into a laptop and then begins moving laterally, understanding what happened before that movement provides important context for determining whether it is legitimate activity or part of an attack.

Speed is equally important. Adversaries are increasingly using automation and AI to accelerate their activities, while enterprises are generating terabytes and petabytes of data. Humans alone cannot effectively process all that information.

We therefore need AI designed to deliver defensive outcomes, and that capability needs to be integrated into the security platform itself.

CIO&Leader: As AI agents begin interacting with other AI agents across partner ecosystems, what should CISOs build today to ensure long-term identity governance?

FABIO FRATUCELLO: The fundamental identity concepts remain similar even when interactions extend beyond a single organisation.

Businesses have long connected with third parties through SaaS platforms. They use identities to authenticate, exchange data and process information.

The medium is now changing. Instead of a person logging into a browser, an AI agent may use an MCP server to communicate with another AI agent. There may also be an AI gateway, API calls or other technical pathways through which information moves.

But the control capabilities remain broadly similar.

When CrowdStrike thinks about AIDR, we are therefore not looking only at compute. We also look at identity interactions across the broader ecosystem.

Organisations need to understand whether there will be MCP calls, API calls or AI gateways and determine what protection mechanisms apply at each level.

The principles of visibility, detection and response remain central. The technology pathways may change, but the need to understand identities, monitor their behaviour and respond when something goes wrong remains fundamentally the same. ■



ET Edge INSIGHTS

CIO&LEADER studiotalks

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

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

Follow us: @CIOandLeader

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 | KARJAT

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

INDIA'S LEADING B2B INTELLIGENCE, COMMUNITY & MARKET PLATFORM