CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

MAY 2026 • VOLUME 15 • ISSUE 02 • ₹150

SCALING AI WITH DISCIPLINE

Vijay Balakrishnan, CDIO of Godrej Enterprises Group, on building enterprise-wide AI through intelligent platforms, AI agents, and responsible AI governance. PG. 10

Responsible Data Wins

Manoj Kern
CIO, Prudent Insurance Brokers
PG. 25

Rise of Digital Workers

Arun Shetty
CTO, Cisco India & South Asia PG. 33

A 9.9 GROUP PUBLICATION
cioandleader.com
cioandleader/
cioandleader

CIO&LEADER studiotalks

CIO&LEADER STUDIOTALKS—
WHERE TECHNOLOGY MEETS THE SPOTLIGHT!

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

WHY JOIN STUDIOTALKS?

- Engage in powerful conversations that shape the future of enterprise IT.
- Share your expertise in a high-impact, TV-style format.
- Be featured among India's top technology leaders.

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

SECURE YOUR SPOT NOW!

For more information
Jatinder Singh
Chief Editor, Enterprise Tech Publications
ET Edge - The Times Group
jatinder.singh1@timesgroup.com, +91 9718154231

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

Follow us: @CIOandLeader

Could OpenAI's services push signal more trouble for Indian IT?

OpenAI's decision to launch DeployCo marks more than just another expansion initiative. It signals a deeper shift in the global technology services landscape, one that could intensify the pressure already building on India's $250-billion IT services industry.

By moving beyond AI models into enterprise implementation and on-ground deployment, OpenAI is entering territory long dominated by Indian IT firms. Through DeployCo, the company will directly help enterprises integrate AI into business operations, workflows, and mission-critical systems. In effect, the AI pioneer is stepping into the services layer itself.

For decades, Indian IT majors such as Tata Consultancy Services, Infosys, Wipro, and HCLTech built globally successful businesses around manpower-led outsourcing, application development, infrastructure management, ERP implementation, testing, and long-term support contracts. Revenue growth was closely tied to the scale of engineering talent deployed across projects.

AI is beginning to disrupt that equation. Tasks that once required large teams can now be automated through generative AI, coding copilots, intelligent workflows, and autonomous AI agents. Enterprises are not really cutting technology budgets, but the focus is swiftly moving toward GPU infrastructure, AI software, analytics, and data foundations rather than traditional services delivery.

This makes the current transition different from earlier downturns. Unlike the cyclical slowdowns witnessed during the 2008 financial crisis, the present wave reflects a structural shift toward automation-first operating models.

Yet disruption does not necessarily indicate decline. Indian IT firms have time and again reinvented themselves through Y2K, cloud, and digital transformation cycles. The next opportunity may lie in AI consulting, governance, systems integration, managed AI operations, and industry-specific AI solutions. The real challenge is speed: how quickly can traditional IT services firms adapt before the rules of the industry are rewritten? ■

"OpenAI is entering territory long dominated by Indian IT firms. It would be interesting to see how quickly traditional IT services firms can adapt before the rules of the industry are rewritten."

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

CONTENT MAY 2026

COVER STORY

10-15

Scaling AI with Discipline

Vijay Balakrishnan, CDIO of Godrej Enterprises Group, on building enterprise-wide AI through intelligent platforms, AI agents, and responsible AI governance— while keeping humans at the center.

Cover Design by: **Shokeen Saifi**

Please Recycle this Magazine and Remove Inserts before Recycling

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**

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**
Regional Sales Head - North: **Sourabh Dixit**
Senior Sales Manager - South: **Aanchal Gupta**

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

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

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

For editorial queries write to:
editor@cioandleader.com

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

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

Editor: **Vikas Gupta**

9.9 GROUP

CIO MOVEMENTS

Pankaj Gupta Appointed Chief Digital Officer at Ujjivan Small Finance Bank

Pankaj Gupta moves to Ujjivan Small Finance Bank from Karnataka Bank to lead digital strategy, analytics-driven transformation, customer experience, and scalable technology-led growth initiatives.

Devesh Verma Appointed Chief Digital Officer at Bank of Maharashtra

Devesh Verma joins Bank of Maharashtra from Poonawalla Fincorp to drive digital banking transformation, innovation, customer-centric platforms, and enterprise technology modernization initiatives.

Omprakash Khawse Appointed Head – IT Infrastructure at Aditya Birla Money

Omprakash Khawse takes on a new role at Aditya Birla Money after moving from Bajaj Finserv Asset Management to strengthen IT infrastructure strategy, operational resilience, and tech modernization initiatives.

Ashok Mysore Appointed CEO at Global Green Grid Data Centers

Ashok Mysore joins Global Green Grid Data Centers from CtrlS Datacenters to scale sustainable, AI-ready digital infrastructure and next-generation data center ecosystems.

Mandar Ghatnekar Appointed Chief Technology Officer at Biocon

Mandar Ghatnekar takes on the CTO role at Biocon after serving at Biocon Biologics, driving digital innovation, AI adoption, and technology transformation initiatives.

Preeti Singh Appointed Chief Information Security Officer at Teciem

Preeti Singh joins Teciem from OSTTRA to lead cybersecurity strategy, governance, risk management, compliance frameworks, and secure enterprise transformation initiatives.

Rajesh Mittal Appointed Chief Information Officer at Sterling and Wilson

Rajesh Mittal moves to Sterling and Wilson from Altius Telecom Infrastructure to lead enterprise IT strategy, digital transformation, and technology-driven operational excellence initiatives.

Samar Gupta Appointed Group CIO & Senior Vice President at Spark Minda

Samar Gupta joins Spark Minda from Anand Group to drive IT-led business transformation, enterprise technology strategy, and data-driven innovation initiatives.

Tushar Zade Appointed Chief Transformation Officer at Granules India

Tushar Zade joins Granules India from Aurigene Pharmaceutical Services to lead AI-led transformation, process excellence, and enterprise-wide digital innovation initiatives.

Raghav Grandhi Appointed Chief Information Security Officer at Ramoji Group

Raghav Grandhi joins Ramoji Group from Spandana Sphoorty Financial to strengthen cybersecurity strategy, enterprise risk management, compliance, and digital security governance frameworks.

Raman Pillai Appointed Chief Digital & AI Officer at SISL Infotech

Raman Pillai takes on a new mandate at SISL Infotech after moving from VerSe Innovation to spearhead AI strategy, cloud transformation, scalable digital platforms, and enterprise modernization initiatives.

Kunal Handa Appointed Chief Information Officer at Greenply Industries

Kunal Handa joins Greenply Industries from Eureka Forbes to lead digital strategy, SAP transformation, enterprise IT governance, and manufacturing technology modernization initiatives.

Swapnasarit Singh Appointed Chief Information Officer at Moon Beverages

Swapnasarit Singh joins Moon Beverages from BCH Electric to lead enterprise IT strategy, SAP transformation, infrastructure modernization, and digital innovation initiatives.

Ruma Kishore Appointed Chief Digital & Information Officer at Titan Company

Ruma Kishore embarks on a new leadership role at Titan Company after a stint at Unilever to accelerate digital transformation, enterprise IT strategy, customer experience innovation, and data-driven growth initiatives.

Kiran Mani Appointed Managing Director, APAC at OpenAI

Kiran Mani joins OpenAI from Jiostar to lead regional growth, strategic partnerships, and enterprise AI adoption initiatives across the Asia-Pacific market.

Dr. A Shiju Rawther Appointed CITO at CareEdge Group

Dr. A Shiju Rawther joins CareEdge Group from SBI Mutual Fund to lead enterprise technology strategy, digital innovation, and transformation initiatives across the organization.

ChatGPT just replaced GPT-5.3 with GPT-5.5 — Here's what actually changed

OpenAI recently released GPT- 5.5 Instant and made it the new default model for ChatGPT, replacing GPT-5.3 Instant. No pop-up alert. No countdown timer. Just a switch.

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

OpenAI recently released GPT- 5.5 Instant and made it the new default model for ChatGPT, replacing GPT-5.3 Instant. No pop-up alert. No countdown timer. Just a switch.

GPT-5.5 Instant is not a completely new model family. It sits within the GPT-5 generation, a series OpenAI began rolling out in mid 2025. It appears to be an iterative evolution within the GPT-5 series rather than a wholly separate model family.

The new model is rolling out to all ChatGPT users, replacing GPT-5.3 Instant as the default. The clearest improvement is in accuracy. Reportedly, in internal evaluations, GPT-5.5 Instant produced 52.5% fewer hallucinated claims that GPT-5.3 Instant on high-stakes prompts covering areas like medicine, law, and finance. And, that is a significant enhancement for users who rely on ChatGPT for professional research or factual lookups.

On the AIME 2025 math test, the new model scored 81.2, compared to 65.4 for the older model. According to reports, it also outperformed its predecessor on the MMMU-Pro multimodal reasoning benchmark, scoring 76 versus 69.2.

OpenAI has replaced GPT-5.3 with GPT-5.5 Instant, boosting accuracy, reasoning, and memory features.

The personalised push

Beyond accuracy, OpenAI has made enhancements in memory and personalisation with this update.

GPT-5.5 Instant can now use its search tool to refer back to past conversations, files, and Gmail to give more personalised and curated answers. These features are available to Plus and Pro users on the web, with plans to roll it out to mobile soon.

It matters to users because most AI chatbots treat every conversation like it's the first one. This new model aims to break that pattern. If you told ChatGPT last month that you live in Mumbai or any other particular place, the model can now factor that in when you ask for restaurants, places to visit kind of suggestions, without you repeating yourself.

OpenAI is also introducing memory sources across all ChatGPT models, giving users visibility over what context was used to personalise responses. Users can delete outdated sources or correct them if the answer was wrong, and memory sources are not shown to others if a chat is shared.

What happened to GPT-5.3?

GPT-5.3 does not disappear immediately, but it is on a clock. For paid users, GPT-5.3 Instant remains available for three months, accessible through model configuration settings, before being retired.

For developers, the GPT-5.5 model is available

through the API as "chat-latest", with GPT-5.3 available as an option for paid users for only three months.

Free users do not have such option. When OpenAI changes the default, free users are going to get what they are given.

How did ChatGPT get here?

To understand this update, it helps to trace the GPT-5 family's evolution over the past year.

OpenAI launched GPT-5 in August 2025, according to reports GPT-5's responses with web search enabled were around 45% less likely to contain a factual error that GPT-4o. With thinking enabled, that figure rose to approximately 80% fewer factual errors than OpenAI o3.

GPT-5.1 followed in November 2025, improving response speed and reasoning. GPT-5.1 Instant introduced adaptive reasoning, deciding when to think before responding to more challenging questions, resulting in more thorough answers while still responding quickly.

GPT-5.2 arrived in December 2025, focusing on reliability and reducing errors in complex domains. GPT-5.3 followed in early 2026 and was itself an improvement in conversational tone and web search quality. GPT-5.3 Instant focused on reducing unnecessary dead ends, caveats, and overly declarative phrasing that can interrupt the flow of conversation.

GPT-5.4 then arrived with a focus on coding and agentic tasks. GPT-5.4 brought together advances in reasoning, coding, and agentic workflows, incorporating the coding capabilities of GPT-5.3-Codex while improving how the model works across software environments and professional tasks.

Now GPT-5.5 Instant takes the default slot.

GPT-5.5 beyond "Instant"

There is a separate, more powerful version of GPT-5.5 that was released weeks before the Instant upgrade.

GPT-5.5 rolled out to Plus, Pro, Business, and Enterprise users in ChatGPT and Codex, with GPT-5.5 Pro also rolling out to Pro, Business, and Enterprise users in ChatGPT.

The benchmark numbers for GPT-5.5 proper are significant. On GDPval, which tests agents' abilities to produce well-specified knowledge work across 44 occupations, GPT-5.5 scores 84.9%. On OSWorld-Verified, which measures whether a model can operate real computer environments on its own, it reaches 78.7%.

The practical use cases are real. Derya Unutmaz, an immunology professor at the Jackson

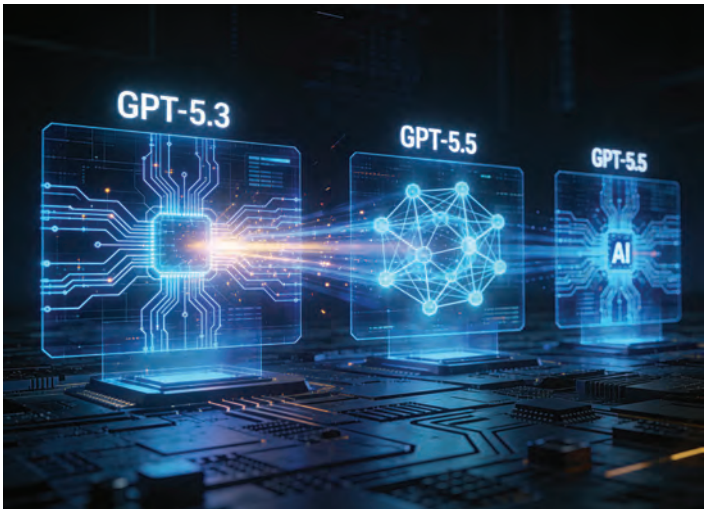

GPT-5.5 Instant becomes ChatGPT's default model, offering better accuracy and personalization.

Laboratory for Genomic Medicine used GPT-5.5 Pro to analyse a gene-expression dataset with 62 samples and nearly 28,000 genes, producing a detailed research report that he said would have taken his team months.

It signals where OpenAI is pushing the product: not just chat, but deep professional work.

Does the backlash problem still exist?

OpenAI has a pattern worth noting. It retires old models, users push back, and the company sometimes adjusts. When OpenAI withdrew its GPT-4o model, there was significant backlash from users who related to the model's personality. Despite the outcry, GPT-4o was deprecated in February 2026.

The emotional attachment users form with specific model versions is a real design problem. People don't just use these tools — some genuinely feel they "know" a particular model's style, tone, and quirks. When that changes overnight, it feels personal.

OpenAI addressed this directly when GPT-5.1 launched. The company said that going forward, when new ChatGPT models are introduced, the approach is to give people ample space to evaluate what's changed and share feedback, and that sunset periods will be communicated clearly and with plenty of advance notice.

Whether GPT-5.5 Instant's rollout reflects that promise is debatable. Users got the switch on May 5. The three-month window for paid users to revert to GPT-5.3 exists, but free users don't get that grace period at all. ■

Cognizant plans to cut upto 15,000 jobs under project LEAP— Will India bear the cost?

Cognizant recently posted its first quarter earnings and buried a significant announcement inside the financial disclosures. The IT services company introduced Project LEAP, a restructuring program that, based on the numbers it disclosed.

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

Cognizant said it expects to record costs of US $230 million to US $320 million in connection with project Leap, with substantially all of the costs expected to be incurred in 2026. That figure consists of US $200 million to US $270 million in employee severance and other personnel- related costs, and US $30 million to US $50 million in other charges.

The company did not confirm a specific headcount number, but the mathematics did it all.

What the numbers suggest

Cognizant has not officially stated how many employees will lose their jobs. According to multiple reports, the estimated scale of layoffs is based on projected severance payouts and average employee compensation levels across regions. Using an estimated average salary in India, it is being estimated that around 12,000 to 13,000 India-based employees could be affected.

With global headcount closer to 357,600

employees, a reduction of 15,000 represents close to 4-5% of Cognizant's total headcount. The number is large enough to reshape the organization but not large enough to operationally destabilize it, which is precisely why fits the profile of a managed restructuring rather than a crisis response.

Higher compensation levels in markets such as the US would mean fewer employees affected for the same cost pool. In practical terms: the same severance budget that eliminates 12,000 to 13,000 roles in India might only affect a few hundred in the United States. The arithmetic of global labour arbitrage, the very model that built Cognizant, is now the same mechanism determining who gets cut.

What is Project LEAP?

Cognizant introduced Project Leap alongside its Q1 2026 results, saying it expects the program to generate US $200 to US $300 million in savings in 2026 and drive full-year adjusted operating margin guidance to 16.0% to 16.2%.

CEO Ravi Kumar S has framed Project Leap as a structural shift, not a cost exercise. Kumar stated that the company is on the journey to get to the operating model, with the goal of reshaping the talent pyramid as digital labour, software and AI, starts handling jobs that once went to people.

Kumar also disclosed that the company has more than 5,000 AI engagements underway and that nearly 40 percent of code is now being AI-assisted. That figure is notable. It means automation is already doing measurable work across Cognizant's delivery operations. Project Leap formalises the logical next step: shrinking the human workforce to match.

CFO Jatin Dalal told analysts that Project Leap is about driving cost savings through a "cost of delivery" model — fewer humans per unit of work.

The roles at risk

Project Leap targets roles in application maintenance, business process outsourcing, and traditional IT support ; functions that automation tools have increasingly taken over.

These are not niche positions. They represent the bulk of entry and mid-level employment at most Indian IT services firms. CEO Ravi Kumar S aims to build a "broader and shorter pyramid," with fewer entry-level positions, fewer middle-level positions, and more digital tools and AI-driven delivery.

Engineers with six to twelve years of experience in routine roles like manual testing, basic application support, traditional database administration, and standard development work, may face the highest risk. These are precisely the roles that AI tools can either fully automate or augment to the point where one engineer does the work of three.

Project LEAP is set to cut costs and reshape its workforce as AI adoption reduces demand for traditional IT roles.

The irony is pointed. For two decades, India's IT sector built its global dominance on exactly those roles. Now those roles are the first on the list.

India – where the impact might land

Cognizant employs over 250,000 workers in India. India accounts for approximately 72 percent of the company's global workforce. That concentration is why the country faces a disproportionate share of the restructuring impact.

A substantial portion of the job cuts is expected to be concentrated in India, where the company has its largest workforce and where cost structures allow for broader adjustments. Cities like Bengaluru, Hyderabad, Pune, and Chennai — which built entire ecosystems around IT employment — might feel the impact directly.

There is also a timing complication. Cognizant simultaneously announced it will hire more than 20,000 freshers in 2026. That appears contradictory at first. But it is not. Technically if we observe the industry pay scales, freshers cost less, carry no legacy habits, and can be trained directly on AI-first workflows. The company is not abandoning hiring, but it is restructuring who it hires and at what level in the pyramid.

A pattern across the sector

Cognizant is not doing this alone. Tata Consultancy Services laid off approximately 12,000 mid-level and senior managers in July 2025, citing essentially the same reason: workforce mismatch with new technology demands. Globally, more than 40,000 tech sector employees lost their jobs to AI-driven restructuring in April 2026 alone.

Freshworks revealed an 11% workforce reduction plan due to AI product and engineering changes. Atlassian announced a reduction in its employee base by 10%. Coinbase announced layoffs accounting for 14% of its workforce.

This is not a cycle. It is a structural shift. The workforce models that sustained India's IT industry for twenty years — large pyramids of developers, testers, and support staff — are being replaced by leaner teams augmented by AI tools. ■

SCALING AI WITH DISCIPLINE

Vijay Balakrishnan, CDIO of Godrej Enterprises Group, on building enterprise-wide AI through intelligent platforms, AI agents, and responsible AI governance— while keeping humans at the center.

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

Today, large enterprises are no longer struggling to identify AI use cases. The real challenge lies in deciding where to begin, scaling AI responsibly, and aligning people, processes, and platforms behind a unified transformation strategy.

Godrej Enterprises Group (GEG) is driving this transformation at significant scale. A diversified Indian conglomerate spanning both consumer and industrial sectors, the group operates primarily through Godrej & Boyce and its affiliated businesses. Its portfolio ranges from home appliances and security solutions to advanced engineering systems supporting India's space, infrastructure, and industrial programs.

With operations across 14 business verticals, a presence in more than 60 countries, and annual revenues nearing Rs 21,000 crore, the group has accelerated modernization initiatives across manufacturing, supply chain, customer experience, finance, procurement, and commercial operations over the past year.

Leading this transformation is Vijay Balakrishnan, who brings more than two decades of experience in driving digital transformation within manufacturing-led enterprises. Under his leadership, GEG has moved beyond foundational cloud modernization and platform consolidation to build enterprise-wide AI capabilities through initiatives such as Factory 360 and Amethyst, its internal AI and intelligence orchestration platform.

In a recent interaction with Jatinder Singh, Chief Editor, ET Edge CIO&Leader, Balakrishnan discusses how GEG is scaling AI across the enterprise, why responsible AI must be embedded from the outset, how AI agents are being integrated into business functions, and why the future of enterprise AI will depend less on experimentation and more on disciplined prioritization. Excerpts from the interaction:

CIO&Leader: GEG has been accelerating its digital and AI transformation journey over the last few years. Could you walk us through how that journey evolved, starting from foundational modernization and cloud adoption to building enterprise-wide AI capabilities like Amethyst?

VIJAY BALAKRISHNAN: This is the fourth major transformation program I've led in my career, and one thing I've learned is that every successful transformation kicks off with strong foundations. At Godrej Enterprises Group, the main focus was on platform modernization, cloud-first adoption, and the creation of a long-term technology strategy that could scale with the business. When I joined, the organization was already committed to this direction, with a strong emphasis on accelerating and completing the modernization journey.

In many ways, joining at that stage proved advantageous because we were able to learn from industry experiences, follow best practices, and avoid several of the early pitfalls that organizations faced during digital transformation initiatives. Alongside infrastructure and platform modernization, we also started reimagining enterprise processes,

building group-wide systems, and strengthening analytics capabilities across the organization.

Our approach was different. We did not spend years building a huge, centralized data program or pick just one platform partner, then work through AI use cases one at a time. Instead, we focused on being agile and experimenting, running several POCs, pilots, and phased rollouts at the same time. This sped up our progress and cut down timelines by months.

The biggest change came with generative AI. Suddenly, AI was almost plug-and-play. Our systems were already producing reliable data, giving us a solid foundation. However, gen-AI also brought a new challenge: hallucination. Because these models are basic and built outside the company, organizations cannot fully control the information they use.

That is the point where intelligence becomes critical. I often compare it to hiring an experienced professional from outside the company. They may bring deep expertise and broad knowledge, but they still need guidance on how your organization operates, which policies to follow, how decisions are made, and how work should be executed within your environment. In the physical world, that guidance comes from policy documents, managers, peers, and organizational culture. In AI, that guidance must come through an enterprise intelligence layer.

This idea led to the creation of Amethyst.

Amethyst is our enterprise intelligence layer. It helps fit AI into the real needs of our organization. With Amethyst, AI systems and agents use trusted company knowledge, processes, and governance, which lowers the risk of hallucinations and makes them more reliable.

Today, we are close to completing the foundational aspects of our modernization journey. While we had already begun experimenting with AI, the progress over the last 18 months has significantly accelerated our AI roadmap. We are now driving impact across multiple internal processes and have already deployed several production-grade AI agents across the enterprise. More importantly, we continue to rapidly scale these capabilities.

From the beginning, security has remained integral to our approach. As a result, we are now looking at AI from two perspectives simultaneously: cybersecurity for AI — ensuring AI systems remain secure, governed, and trustworthy — and AI for cybersecurity, where AI itself becomes a force multiplier for strengthening enterprise security operations.

CIO&Leader: Many enterprises are still struggling with pilot fatigue, production-scale deployment, and proving real business impact from AI initiatives. What do you think has enabled you to accelerate AI adoption more effectively?

VIJAY BALAKRISHNAN: We learned early on that using AI in isolated areas does not create real business value. For example, AI in customer service is only as good as the supply chain, service, and logistics systems behind it. If those systems are not connected, AI just makes inefficiencies more obvious.

That's why we take a holistic approach. We look at the whole design-to-delivery value chain and use AI across connected processes, not just in separate functions. The main challenge now is not finding where to use AI but figuring out the best order to scale it.

We set our priorities based on business impact, feasibility, data readiness, process maturity, and how well people adopt new tools. By focusing on these areas, we have moved past pilot fatigue and scaled AI more effectively throughout the company.

CIO&Leader: Could you explain the key AI use cases you are focusing on across customer experience, manufacturing, and enterprise operations, and how you decide where AI can create the maximum impact?

VIJAY BALAKRISHNAN: Our AI strategy is fundamentally hybrid and closely aligned with our wider platform vision. On the enterprise side, we have invested heavily in cloud-native ERP, CRM, and planning platforms with robust APIs that enable both data access and workflow automation. For specialized manufacturing needs, we built Factory 360, which today supports over 30 factories, alongside Amethyst, our enterprise intelligence and orchestration layer.

Amethyst, our enterprise intelligence and AI orchestration layer, has three main roles. It acts as the intelligence layer, serves as an agentic AI platform built on frameworks like LangGraph and LangChain, and works as the orchestration engine that connects workflows across systems and functions. We do not view AI as a set of separate use cases. Real impact happens when manufacturing, supply chain, logistics, planning, and customer service all work together as one connected value chain.

One of our largest initiatives has been Customer UID, which unifies sales, service, and marketing interactions inside a single customer platform. By combining transactional data, sentiment analysis, and social insights, we are enabling far more personalized customer engagement.

We are also launching a multilingual Voice AI platform that lets users switch easily between Indian languages during conversations. This delivers customer interactions that feel almost human and ensures every call is answered.

In manufacturing, we use AI for predictive maintenance, sensor-based asset monitoring, full traceability, and improving sustainability in our factories. AI is also being added to connected products and IoT systems, making them smarter and more responsive.

Inside the company, we are automating complex workflows across different functions. For example, in our furniture business, we have automated most of the B2B order processing from start to finish.

> “RESPONSIBLE AI CANNOT BE AN AFTERTHOUGHT — IT HAS TO BE BUILT INTO THE ARCHITECTURE FROM DAY ONE.”

In all these projects, our goal is the same: to improve agility, business efficiency, and customer experience on a large scale. We look at AI agents almost like digital employees. Every agent has a human owner who is ultimately accountable for its performance and outcomes.

CIO&Leader: Have you been able to track any efficiency gains or business outcomes from these deployments?

VIJAY BALAKRISHNAN: It’s a complex question because ROI depends on the use case. I always focus on early indicators for ROI. In the end, ROI appears in the P&L, and for AI projects, it usually affects the bottom line.

Top-line impact is harder to prove. For example, if sales improve, how much credit goes to AI versus everything else happening in the organization? That’s always debatable.

But when it comes to productivity, we are seeing improvements of 10 to 15 percent, and in some areas, more than 50 percent.

Sometimes, it is about increasing output with the same number of people, which helps us avoid extra costs. In the B2B order-booking example, we have increased throughput by nearly five times with the same team.

In some cases, our dependence on external agencies has been reduced dramatically.

But I always return to early indicators. If people focus only on immediate P&L results, they may become skeptical too soon and miss bigger opportunities.

> INSTEAD OF TREATING SHADOW AI AS A THREAT, WE CHOSE TO CREATE A SECURE, FAIL-SAFE ENVIRONMENT THAT ENABLES EMPLOYEES TO INNOVATE RESPONSIBLY.

CIO&Leader: And to what extent does this ecosystem rely on partnerships with companies like Microsoft, OpenAI, or hyperscalers, versus on capabilities built internally?
VIJAY BALAKRISHNAN: Our approach is hybrid. Amethyst itself runs on AWS, our strategic cloud partner. We use some native AWS components and build others ourselves. But the real intelligence layer is custom-built internally. At the same time, we leverage native AI agents within our CRM and ERP platforms through Amethyst.

CIO&Leader: And how critical is proprietary data in making these platforms effective?
VIJAY BALAKRISHNAN: TVery important. Our philosophy is simple: first, make processes lean, then digitize them, and build strong data before scaling AI. In other words, optimize the process, digitize it, build a solid data foundation, and add AI only then.

Earlier, each of these stages required independent focus, and organizations would spend months progressing from one layer to the next. Today, technology has evolved to the point where companies can move across these dimensions simultaneously and progressively, accelerating the overall transformation journey.

One of the strengths of our organization is the way we have structured functional councils across manufacturing, services, and marketing. Since we operate through 14 business units, these councils bring together functional leaders to establish common process ownership and governance. This allows us to standardize processes even before digitization begins, creating a firm foundation for transformation.

The second aspect is where things become more subtle. Traditionally, platform transformation and data transformation were treated separately. However, many modern enterprise platforms now inherently ensure reliable, high-quality, and consistent data through native APIs. As a result, the earlier split focus between platform modernization and data readiness is no longer necessary in many cases.

There are still situations where dedicated data interventions are required, and we continue to address those separately. But for core enterprise systems such as CRM, ERP, S&OP, and HCM, the architecture itself is increasingly designed to deliver structured, reliable data at scale. If these platforms are implemented well — especially after lean process standardization — they naturally generate better-quality data.

And then comes AI. As I mentioned earlier, the activation time for AI has reduced dramatically.

CIO&Leader: What does responsible AI mean in practice for Godrej, and how are you building governance, accountability, explainability, and monitoring frameworks for AI use across the organization?

VIJAY BALAKRISHNAN: The first thing we did was establish an AI Advisory Board comprising top management, business leaders, and key functional heads. Every major AI initiative is reviewed and vetted through that governance structure.

Second, for us, responsible AI fundamentally means keeping humans in the loop. We do not believe in completely autonomous systems because there are too many socioeconomic variables, business nuances, and unpredictable drifts involved.

For us, AI agents work much like digital employees. Each agent has a human owner who is responsible for its performance, decisions, and results. Just as managers are responsible for their teams, there must be clear human accountability for AI systems.

An additional essential pillar for us is explainability. This becomes especially important in areas such as demand or sales forecasting, or in decision-support systems, where outcomes could influence incentives or business decisions. Users must clearly understand why a recommendation is being made.

We have added governance features directly into Amethyst to reduce hallucinations, cut down on bias, and make sure responses fit our business context. The platform also has clear rules: if a question is not related to work, the system will not answer.

Overall, our philosophy is clear: responsible AI should be part of the architecture and operating model from the start. For us, it is truly 'responsible AI by design.'

CIO&Leader: How are you addressing the increasing challenge of Shadow AI, which many CIOs and CISOs see as a major governance and security concern?

VIJAY BALAKRISHNAN: Instead of fearing Shadow AI, we have built a safe environment that encourages innovation. We have added governance, hallucination control, and bias filtering directly into the main platform.

In fact, I think it should be encouraged, as long as it is done responsibly.

We have been investing significantly in employee AI awareness and developing AI capability across the organization. Last year, we provided about 600,000 hours of AI-related training, and close to 6,000 employees completed foundational AI learning programs. The wider objective is to make AI adoption increasingly inclusive and empower employees to use these technologies confidently and responsibly. With a secure platform like Amethyst, employees can confidently create their own knowledge agents, and soon, even their own actionable agents.

To strengthen governance even more, we are building an AI Command Center. We already have an early version running, and by the end of this year, it will be fully set up.

We are also cautious about free public LLMs because of security risks. Employees might accidentally share sensitive information, underscoring the importance of cybersecurity awareness and responsible platform use..

CIO&Leader: Are you seeing a shift in the balance between automation, augmentation, and the kind of skills the organization now prioritizes?

VIJAY BALAKRISHNAN: As a company, we have core values and responsibilities that do not change. We are mainly working to reduce human involvement in agency and contract work. If we do not use AI to improve efficiency in these areas, we are not serving the organization well.

For our internal employees, the focus is on helping them do more with AI. As a growing company, we have been adding to our workforce every year. AI might slow hiring in some areas, but we still need strong engineering talent. That will not change.

People are open to AI when it helps with the frustrating parts of their daily work. If an employee sees that an agent can handle repetitive, boring tasks, they are more likely to use it.

CIO&Leader: With GEG planning to invest nearly Rs 1,200 crore into AI-led initiatives, what are the key strategic priorities and transformation areas you see forming the organization's AI roadmap over the next few years?

VIJAY BALAKRISHNAN: For us, employee safety is still one of our top priorities. We are using AI, computer vision, drones, and Amethyst-powered monitoring systems to create a smarter safety system in our factories and customer areas. The goal is to move from reacting to problems to using proactive and predictive systems that spot risks and help make workplaces safer with real-time monitoring and intelligence. ■

What's holding CIOs back on digital sovereignty?

If there is any doubt in the industry whether digital sovereignty has become a boardroom priority or not, a recent report titled "Navigating Digital Resilience" by SUSE has settle it down.

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

Nearly 98% of IT leaders surveyed by SUSE has said digital sovereignty is a priority. Yet only 52% are actively taking steps to achieve it. For CIOs, this gap between ambition and execution isn't just an operational detail, it is a startetegic vulnerability. As AI adoption accelerates and regulatory expectations intensify, the window to act decisively is narrowing.

AI as a catalyst or complicator

The report highlights the tension that lies in the relationship between AI and sovereignty. On one hand, 64% of IT leaders believe AI transparency; control over model training and provenance, will be the top driver of digital resilience over the next five years. On the other, when offered a hypothetical 20% budget increase, organisations over-

whelmingly prioritise AI implementation over sovereignty investments.

This signals a troubling pattern: the pressure to adopt AI is outpacing efforts to manage the risks it introduces.

What's driving action — and what isn't

The research reveals that external pressure remains the primary catalyst for sovereignty intiatives. A full 41% of respondents admit they only act when required by customers or regulation. While 45% have included sovereignty requirements in recent RFPs and 42% selected vendors on that basis, the reactive posture suggests many enterprises are still waiting for external forces to compel change rather than treating sovereignty as a proactive differentiator.

The report highlights notable regional variation:

- **India** leads with 62% of respondents describing digital sovereignty as a genuine strategic priority they are actively investing in.
- **Germany and Japan** follow at 57% each.
- **The U.S.** stands at 52%, with 61% expressing optimism about digital resilience and 41% already having a formal sovereignty strategy.
- **France** trails at 39% actively investing.

Control as the common thread

While definitions of digital resilience vary, the report finds organisations converging around a core priniciple: control.

Top priorities include:

- Cybersecurity and threat detection (63%)
- Multi-cloud or hybrid diversification (52%)
- Backup and recovery (45%)
- Continuous monitoring (44%)
- Resilience is no longer just about surviving disruptions—it's about maintaining control in increasingly complex, AI-driven environments.

The hyperscaler dilemma

Enterprises remain deeply reliant on hyperscalers, evne as sovereignty concerns grow. A notable 65% of respondents say hyperscalers are relevant for supporting sovereign workloads, creating a delicate balancing act between scale, convenience, and jurisdictional control.

CIOs broadly agree digital sovereignty is critical, but most are still struggling to translate intent into concrete action.

This tension is driving demand for open, interoperable solutions and regional ecosystems that offer flexibility without lock-in.

What CIOs should consider now

- **Audit the gap:** Assess whether your organisation's stated sovereignty priorities are matched by budget allocation and concrete initiatives.
- **Embed sovereignty into AI strategy:** Don't treat governance as a bolt-on. Build control over data, models, and infrastructure into AI programmes from inception.
- **Move from reactive to proactive:** Waiting for regulation or customer pressure is a losing strategy. Early movers will shape vendor relationships and competitive positioning.
- **Evaluate hyperscaler dependency:** Understand where reliance on global providers creates risk—and where open, interoperable alternatives can provide optionality.
- **Tailor strategy to geography:** Recognise that regulatory and market conditions vary. A one-size-fits-all approach to sovereignty will fall short.

Looking forward

The Navigating Digital Resilience report delivers a sobering message: nearly everyone agrees digital sovereignty matters, but far fewer are acting on it. For CIOs, closing this gap is no longer optional. In an era where AI is both the engine of innovation and a source of new risk, control over infrastructure, data, and models is becoming the foundation of competitive resilience.

The question is no longer whether to prioritise sovereignty—it's whether you're moving fast enough to make it real. ■

How QBurst Is turning enterprise AI into real business impact

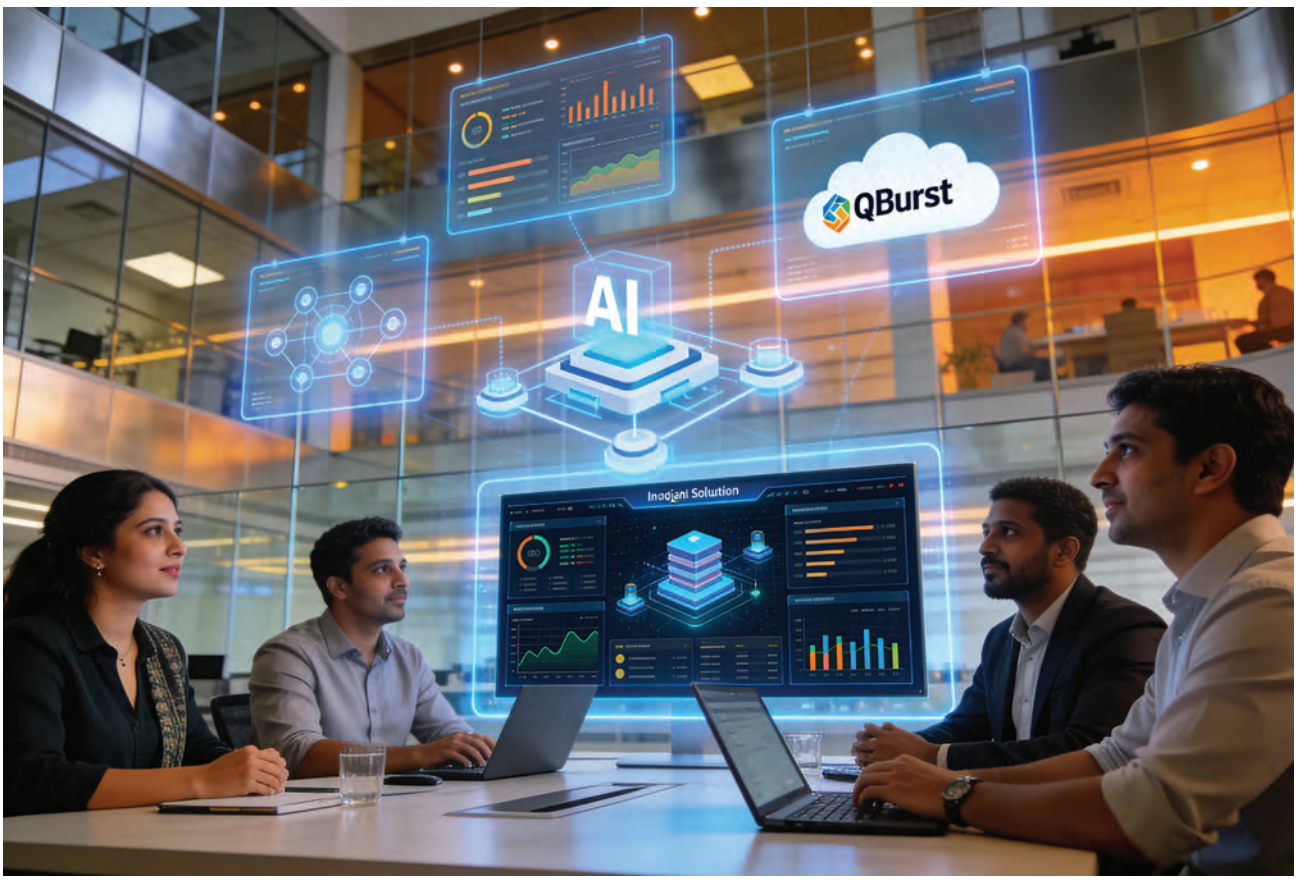

Global technology consulting firm shares its blueprint for driving enterprise AI success

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

Over the past three years, enterprise AI has followed a familiar arc—surging hype, aggressive experimentation, and increasingly uneven outcomes. While boardrooms rushed to embrace generative AI, many initiatives stalled at the pilot stage, unable to translate promise into measurable business value.

IToday, enterprises are confronting a harder reality: AI success is not defined by models or demos, but by execution at scale. The question has shifted from what AI can do to how it can deliver sustained impact.

According to Arun Ramchandran (Rak), CEO of QBurst, a global technology consulting and software development company that helps enterprises build, modernize, and scale digital systems, the

industry is entering a more grounded phase. "We're moving from experimentation to purposeful deployment," he says. "The focus is no longer can we build AI—it's how does it drive real business outcomes."

In his first year as CEO, Rak has repositioned QBurst with an AI-first focus, including the launch of its High AI-Q brand identity. The company has also set up its US headquarters in Palo Alto, signalling its global ambitions. Backed by an investment of about USD 200 million from Multiples Alternate Asset Management, QBurst is accelerating its shift toward a scaled, AI-led platform.

Fixing the foundations

At the heart of enterprise AI's struggle lies a structural issue: legacy infrastructure and fragmented data environments. Most organizations continue to operate on systems that were never designed for AI—resulting in siloed data, limited scalability, and slow integration cycles.

QBurst says that it follows an approach not to replace these systems outright, but to modernize them incrementally.

"Most organizations still run on legacy software," says Ramchandran. "We integrate AI into these environments, helping businesses modernize their technology stack in a practical, scalable way."

A critical part of this effort is data engineering—cleaning, structuring, and preparing enterprise data for AI. Despite heavy investments in data platforms, many enterprises still struggle with unusable, fragmented datasets. QBurst focuses on making this data usable, building the foundation required for AI to deliver consistent outcomes.

Equally important is integration and orchestration—ensuring that AI systems connect seamlessly with enterprise workflows. Without this, AI remains isolated, unable to influence real business processes.

"Building AI agents will become commoditized," Ramchandran notes. "The real differentiation will come from data readiness and orchestration."

Arun Ramchandran (Rak)
CEO, QBurst

From pilots to production

The early wave of AI adoption was dominated by proofs of concept—chatbots, copilots, and automation tools designed to showcase potential. But many of these initiatives failed to scale, largely because organizations focused on building applications while neglecting the underlying engineering.

"Companies invested in the 'middle layer'—AI applications—but overlooked the first mile of data readiness and the last mile of adoption," Ramchandran explains.

QBurst addresses this gap by focusing on end-to-end execution, from data pipelines to system integration to continuous optimization.

In one instance, the company says that it worked with a global fashion retailer to transform a struggling AI chatbot into a production-grade system. Initially plagued by latency, inconsistent responses, and language issues, the solution was failing to deliver value. "The breakthrough didn't come from changing the model," Ramchandran says. "It came from disciplined engineering and integration."

The result was a scalable, high-performing system that delivered measurable business impact—underscoring a key insight: enterprise AI success is less about algorithms and more about operational discipline.

Enabling the Agentic AI phase

As enterprises move beyond pilots, the focus is shifting toward more advanced capabilities, particularly agentic AI. These systems go beyond generating outputs to executing workflows, interacting with applications, and making decisions.

QBurst is helping enterprises adopt these capabilities with a strong emphasis on control and governance. Rather than enabling unrestricted autonomy, the company designs AI systems with clear boundaries, defined roles, and human oversight.

"Think of AI agents as new employees," Ramchandran explains. "You need to define what they can do, where they operate, and how they are supervised."

QBurst positions agentic AI as an overlay to existing enterprise systems, not a replacement. Core platforms such as ERP and IT systems remain intact, while AI transforms how users interact with them, making processes more intelligent, responsive, and efficient.

From experimentation to execution

The enterprise AI landscape is undergoing a critical transition. The era of hype and experimentation is giving way to one of accountability and outcomes. CIOs and boards are asking tougher questions—around cost, scalability, risk, and ROI.

The companies that succeed in this next phase will not be those that build the most advanced models, but those that fix their data foundations, integrate AI into core workflows, and apply it with discipline and intent.

As Ramchandran puts it, "AI is not magic. It's engineering, discipline, and intent. And when those come together, that's when real transformation happens." ■

Unstructured data management has a new job description

Unstructured data is no longer just a storage problem. As AI adoption rises, enterprises must transform fragmented, high-risk data into governed, cost-efficient, AI-ready business assets.

By **Prateek Kansal** | editor.tech@timesgroup.com

For most of the last decade, unstructured data management was a storage problem. Move it, tier it, archive it, forget it. The tools built to handle file and object data were designed around one core objective: keep costs down and keep the lights on. That era is over.

Unstructured data, which includes user documents, multimedia files, logs, emails, research and instrument data and anything else not in a database now represents roughly 80–90% of all enterprise data. Now that enterprises are storing multiple or dozens of petabytes of this data, along with other developments, its needs and requirements have vastly shifted:

- AI, which runs on this data, requires it to be clean, high quality and cataloged. As CXOs demand AI strategies and clear pathways for ROI without incurring risk, unstructured data

has quickly become a critical enterprise asset, and liability.
- Costs to store and back it up escalate every year with 20% or higher annual growth rates. The situation has significantly worsened in 2026 due to the SSD and DRAM shortages and 30-100% price surge from IT infrastructure and hardware vendors.
- Security teams understand the compounding risk of unstructured data as it is unmanaged, highly and easily accessed and shared across teams and geographies.

The software category built to manage this data is evolving rapidly to meet these demands.

What independent unstructured data management actually means

Given the commonplace use of this term across storage and data management vendors today, let's review the category as it has developed in recent years. Independent unstructured data management software operates across storage environments, including on-premises NAS and object stores, cloud storage and edge locations, to deliver analysis, movement and workflows holistically.

These platforms are distinct from the management consoles bundled with your NetApp, Everpure, Dell, Qumulo, VAST or even Amazon S3 buckets. This independence matters: when your unstructured data spans three to four vendor-native tools, you must patch together partial views of a problem that requires one complete picture. You also cannot execute data management policies across your hybrid storage environments with a storage or cloud vendor-centric tool. Plus, storage vendor methods for tiering data to low-cost storage are proprietary, disruptive to users and limit savings and flexibility.

Modern, storage-agnostic platforms can turn unstructured data liability into a cost-efficient, governed, AI-ready asset that is the foundation for organizational success.

Prateek Kansal
Senior VP of Engineering & Operations, Komprise.

Four forces driving the new face of unstructured data management

Cost optimisation in hybrid IT

Few organizations have accurate, deep analytics and visibility into their unstructured data: types, sizes, growth rates, where it lives, departmental trends, access patterns, and what it costs to store and move. This lack of visibility makes it difficult or impossible to right place data as it ages or as the data becomes less valuable to the organization. Duplicate and orphaned data can also be rampant, and deleted, yet too many organizations don't have insights here either. The time is now to get deep visibility for data lifecycle management as the SSD price surge, driven by widespread memory shortages, may only get worse.
- Policy-driven, flexible tiering based on actual access patterns and showback reporting that ties storage costs to business units are standard capabilities in mature platforms.
- This gives IT a credible basis for cost accountability conversations with finance and allows for analytics driven lifecycle management and capacity reclamation.
- File-based tiering, versus block tiering with storage vendors, allows for full file access at the destination and zero rehydration costs when moving tiered data to new storage.
- Data access for tiered data should be simple and transparent, and the solution should never impede hot data performance.

AI data preparation & workflows

Every enterprise AI initiative eventually runs into the same wall: the data isn't ready. Models need clean, labeled, contextually rich training data; unstructured data in most organizations is a sprawling mess of inconsistent metadata, redundant files, and content that nobody has touched in years. The newest generation of unstructured data management tools is attacking this problem directly, with automated metadata enrichment, content classification engines, and governance capabilities to protect sensitive data from ingestion into AI agentic workflows and RAG pipelines.

IT teams that can reliably index, tag, curate and deliver high-quality unstructured data become genuine enablers of AI projects rather than a bottleneck.

IT also needs ways to prevent AI processing and storage waste. This comes down to automated workflows that can curate exactly the right amount of data to AI and no more, along with deleting copies sent to storage for AI processing once the job's complete.

Ransomware and cybersecurity resilience

Unstructured file stores are among the most attractive ransomware targets in the enterprise. They're massive, often inconsistently monitored, and frequently connected to systems across the organization. Recovery from a ransomware event hitting unstructured data has historically been slow, expensive, and incomplete.

Modern unstructured data management platforms are building

Unstructured data is shifting from a storage problem to a strategic AI, security, and governance priority for enterprises.

security capabilities directly into the data layer:

- Sensitive data detection and mitigation;
- Behavioral anomaly detection on file access patterns;
- Tighter integration with immutable snapshot and air-gap technologies, and;
- Ransomware defense by tiering cold data to immutable object storage that attackers can't touch. This can shrink the attack surface by up to 80%.
- Unstructured data governance and compliance

The compliance surface area for unstructured data has expanded significantly. GDPR, HIPAA, CMMC, NIST, SOC 2, the EU AI Act, and a growing body of state-level data privacy regulation all touch unstructured content in ways that legacy storage tools were never designed to address.

Watertight compliance programs require policy-based automated retention and deletion, sensitive data discovery that can identify PII and regulated content at scale, full audit trails tied to content classification, and role-based permissions that follow data across environments.

What IT managers need to watch

The capabilities described above represent genuine progress, but they also raise the bar for the teams deploying them. These are no longer tools that a storage administrator configures once and monitors occasionally. Teams must think beyond traditional storage metrics and capacity planning to data policy, classification logic and cross-functional governance.

There is also an organizational dimension that IT leaders would be unwise to underestimate. Decisions about data retention, AI readiness, and compliance posture now intersect directly with legal, security, and finance. IT teams that position themselves as strategic partners in those conversations will have influence over outcomes, processes and resources.

Addressing unstructured data management objections

Enterprise adoption of new data management solutions can stall around a predictable set of objections. They're worth addressing directly.

- **"We already have tools from our storage vendors so why add another layer?"** Storage vendor tools are optimized for their own platforms. In a hybrid environment where data moves across multiple systems and clouds, they provide partial visibility at best, and tiering capabilities that don't save enough.
- **"We don't have the budget or headcount for this right now."** The cost of a ransomware recovery event, a failed compliance audit, 2X costs for new Flash storage or a delayed AI project due to data readiness failures is increasingly quantifiable. Build a business case with concrete numbers from your own environment to tell the story.

Seizing the unstructured data business opportunity

The teams that will navigate the next three to five years of enterprise IT most effectively are those that stop treating unstructured data management as a housekeeping function and start treating it as the lever to create new organizational value with AI. The practical question for any IT leader is this: does your current toolset give you genuine visibility and control across all four of these dimensions? If it doesn't, the gap between where you are and where your organization needs you to be is your roadmap. ■

AI in hiring: Balancing efficiency and authenticity in leadership recruitment

AI is transforming hiring, but leadership recruitment still depends on human judgment to assess authenticity, context, and long-term potential

By **Sahil Thakur** | editor.tech@timesgroup.com

The way organisations hire is changing fast. Recent industry analysis suggest that a growing share of employers now use AI in some stage of hiring, particularly for sourcing, screening, and communication, with adoption in HR and recruitment tasks increasingly at a fast pace. AI assists in faster screening, easier access to specialised talent, & the ability to handle large volumes of applications with limited manual effort.

But as hiring becomes more efficient, an important question emerges: is the process also becoming less authentic, especially at the leadership and C-suite level?

Where AI really helps

In tech-driven and product-led businesses, the demand for specialised talent is constant. Roles such as machine learning, data engineering, cyber

security and digital infrastructure require very specific skills that are often time-consuming to find.

Here, AI tools can add real value:

- They scan thousands of profiles in minutes instead of days
- They recognise related skills and experience, even when candidates use different terminology
- They help recruiters build stronger initial pipelines and reduce time-to-hire

For high-volume or highly specialised searches, AI brings structure and predictability. It can match core competencies to job requirements at a scale no human team can manage manually, allowing recruiters and hiring managers to focus their energies on influencing & assessments.

The "AI vs AI" paradox

However, candidates are no longer passive in this transformation. Many now use generative AI tools to refine resumes, tailor cover letters and even draft responses for assessments and emails. Employers and jobseekers are increasingly using AI in their hiring interactions, from automated screening on the employer side to AI-assisted CV and content creation on the candidate side. This creates a paradox; an AI-optimised application goes into an AI-powered applicant tracking system where both ends are relying on patterns, templates, and keyword logic.

In such a system, telling the difference between genuine depth of experience and a carefully engineered profile becomes harder. Over time, candidate profiles can start to look and sound similar. If organisations rely too heavily on automated filters, they risk missing out on talent that genuinely fits the role.

Why leadership hiring is different

This tension is most visible in leadership and C-suite hiring. Executive roles differ significantly from volume hiring for junior or mid-level positions.

Sahil Thakur
Director & Head of Tech-Enabled Businesses, Grassik Search

For leadership search, AI can be used to aide specific parts of the hiring process:

- Verifying basic credentials and career progression
- Mapping the market for people with specific role titles, team sizes or industry exposure

However, it struggles with the softer aspects of a candidature like behavioural aspects, how they react to a particular situation, are they able to inspire teams or do they purely manage teams? among other things.

Gauging behavioural aspects requires a nuanced understanding of not just what happened, but how and why it happened, as well as the trade-offs made along the way. Current AI tools are not yet capable of reliably assessing these elements in a way that can replace human evaluation.

There is also a reputational risk. Over-reliance on automated filters at senior levels can result in unconventional but high-potential leaders being filtered out early simply because they did not follow a linear path or use the "right" language. The very people who may help an organisation navigate disruption may never make it onto the shortlist.

Blending technology with human judgement

Because of these limitations, the most effective organisations are designing hybrid models that blend the strengths of both. While a large majority now use AI for at least one part of their hiring process, most still view human decision-making as critical for final selection, especially for senior to leadership roles.

In leadership hiring, human judgement should remain the final filter. Once AI has helped identify a relevant long list, experienced search professionals and business leaders, need to:

- Explore the story behind the CV, including context, failures, and key decisions
- Test strategic thinking and problem-solving in live, unscripted conversations
- Assess culture fit, values and leadership style over multiple interactions and reference checks

Transparency matters as well. Keeping authenticity at the centre

AI will continue to reshape hiring in the coming years. Adoption is growing, tools are becoming more sophisticated, and the pressure to move faster is not going away. But in leadership recruitment, the real advantage will belong to organisations that understand where automation adds value and where human insight & the capability to bring relevant talent on the table are irreplaceable.

When AI is used to handle scale, clean up information and remove friction from the early stages of hiring, it creates the space for humans to do what only they can do: read context, understand character and judge long-term potential. Organisations that strike this balance will build leadership teams that are not just high-performing on paper, but grounded in real values and credibility. In executive search, authenticity is more than a talking point, it is the basis for trust, resilience and sustainable performance in an unpredictable world. ■

Deploying AI without aligning data security creates asymmetric risk

Manoj Kern, CIO at Prudent Insurance Brokers outlines how insurance brokers must shift to architecture-led data governance in the DPDP era, balancing dual fiduciary-processor roles, dynamic consent, third-party risk, and AI governance to build trust and accountability at scale.

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

As India's insurance broking industry adapts to the DPDP era, the real challenge is no longer compliance in isolation, but building end-to-end data accountability across a complex, multi-stakeholder ecosystem. In this interaction with CIO&Leader, Manoj Kern, CIO at Prudent Insurance Brokers, unpacks how firms must navigate dual fiduciary-processor roles, embed consent into system design, strengthen third-party governance, and address emerging AI risks. His perspective highlights a clear shift from policy-led compliance to architecture-led governance—where trust, accountability, and technology must work in tandem to define the future of insurance advisory. Excerpts from the interaction.

CIO&Leader: As an intermediary between customers and insurers, how do you interpret your role under DPDP — Data Fiduciary or Processor — and where does accountability lie? Additionally, how should insurance brokers rethink data governance across distribution channels and third-party partners?

MANOJ KERN: The honest answer is that an insurance broker occupies both roles and failing to recognise that duality is where most governance gaps originate.

When a customer engages with us to explore, compare, or purchase a policy, we are clearly the Data Fiduciary, determining why data is collected, how it is used, and where it flows. Once the client appoints us and data moves to the insurer for underwriting, issuance, or claims, the dynamic shifts, we act as processors on the client's instructions, while the insurer assumes its own fiduciary responsibilities.

This makes accountability inherently distributed. It cannot sit with a single team, it must be contractually defined and operationally enforced at every point where data changes hands.

To strengthen data governance, I see three urgent priorities.

- **First, the distribution layer.** The highest data risk lies not in technology, but in human behaviour at the point of collection across digital channels, branches, and third-party advisors. If data is over collected or misused, accountability traces back to the broker. Governance must therefore move beyond policies and be embedded directly into tools, workflows, and onboarding.
- **Second, the partner ecosystem.** Insurers, TPAs, and technology vendors all handle customer data. Legacy contracts are no longer sufficient, agreements must clearly define data usage, retention, deletion, and breach obligations. Under DPDP, accountability extends across the entire data chain.
- **Third, consent architecture.** A single interaction can trigger multiple downstream data flows, yet most firms still rely on one-time consent. This must evolve into a dynamic, auditable consent framework integrated with core systems, not as a compliance goal, but a regulatory necessity.

At a broader level, DPDP demands far greater

"Customers are increasingly aware of their data rights, and firms that demonstrate responsible data practices will build stronger trust."

precision in data accountability, something previously seen mainly in frameworks like ISO 27701. For firms without mature privacy practices, the complexity of the broking model makes compliance especially challenging without a deliberate, architectural approach.

Ultimately, accountability rests with leadership. Organisational complexity is not a defence and should not be treated as one.

CIO&Leader: With DPDP requiring granular, purpose-specific consent, how would you opine that managing consent and data sharing when the same customer data flows to multiple insurers?

MANOJ KERN: This is arguably the most operationally complex challenge that DPDP presents to insurance brokers specifically. A customer who approaches an insurance broker for health cover expects the broker to find the best option across the market. That inherently means their data, health history, age, occupation, and existing conditions, needs to travel to multiple insurers for quotation and underwriting. Under DPDP, each of those flows is a distinct processing activity, potentially requiring distinct consent. Managing that without creating an unacceptable consent burden on the customer, while remaining fully compliant, is the design challenge the industry must solve.

The answer, in my view, lies in three interconnected principles.

- **The first is consent layering.** Rather than presenting customers with a single blanket consent at the point of engagement, firms need to architect consent in layers, a foundational consent for the broking relationship itself, and purpose-specific consents triggered as the customer journey progresses. When a customer asks for health insurance quotes, consent for sharing their data with shortlisted insurers should be sought at that moment, clearly articulating which insurers, for what purpose, and for how long. This keeps consent meaningful rather than mechanical.
- **The second is a dynamic consent record.** The days of signed proposal forms being the end of the consent story are over. What DPDP demands, and what good practice requires, is a consent record that is timestamped, auditable, and travels with the data. If a customer later withdraws consent for a specific insurer or purpose, that withdrawal must be honoured and traceable across the entire data chain. This requires investment in a consent management layer integrated with core systems, not bolted on as an afterthought.
- **The third is insurer accountability alignment.** When customer data leaves the broking firm and reaches an insurer, consent obligations do not simply transfer, they must be contractually mirrored. Data sharing agreements must explicitly state what data was shared, under what consent, for what purposes, and what the insurer's obligations are if consent is withdrawn or modified.

There is also a broader industry conversation needed around standardisation. Today, every firm is building its own interpretation of granular consent. IRDAI and industry bodies have an opportunity to define templates and protocols that create consistency across the ecosystem, reducing the risk of consent fatigue.

CIO&Leader: With increasing digitisation, how can brokers embed data minimisation and purpose limitation into platform design from day one?

MANOJ KERN: Data minimisation and purpose limitation are not features you can retrofit into a platform after it has been built. By the time architecture, data flows, and integrations are established, redesigning for privacy becomes exponentially more complex. These principles must be treated as foundational design constraints from the outset.

For insurance broking platforms, three areas are critical.

- **The first is purposeful data collection within regulatory constraints.** This means collecting only what is necessary for a

defined purpose, ensuring underwriting data is not repurposed without consent, and avoiding unnecessary data accumulation.

- **The second is integration architecture with clarity of purpose.** Not all integrations involve personal data; some serve operational functions. Where personal data is involved, APIs must be tightly scoped, sharing only required fields with proper access controls and audit logging.
- **The third is consent-aware architecture.** If data is collected for a specific purpose under consent, the system must enforce that limitation and prevent reuse without fresh consent.

The broader point is that privacy by design is not a regulatory burden, it is a competitive advantage. Customers are increasingly aware of their data rights, and firms that demonstrate responsible data practices will build stronger trust. In a relationship-driven business-like insurance broking, that trust is a critical asset. .

CIO&Leader: With insurers, TPAs, surveyors, and tech vendors in the ecosystem, how should insurance brokers strengthen third-party data governance for end-to-end DPDP compliance?

MANOJ KERN: The broking ecosystem is both an operational strength and a data governance vulnerability. Customer data, often sensitive, flows across multiple entities, and under DPDP, accountability ultimately traces back to the originating firm.

- **The first principle is to treat certifications like ISO 27001 as a baseline,** not a guarantee. Firms must assess how vendors handle data, their breach history, subprocessors, and their ability to honour deletion and consent withdrawal.
- **Second is contractual precision.** Data Processing Agreements must clearly define what data is shared, for what purpose, retention timelines, deletion obligations, and breach notification requirements. Any ambiguity is a governance gap.
- **Third is risk-based tiering.** Governance should be proportionate, with stricter oversight for high-risk vendors.
- **Fourth is managing fourth-party risk.** When vendors use subprocessors, data moves beyond direct visibility, and firms must ensure transparency and control.
- **Finally, governance must be continuous,** with periodic reassessments, audit rights, and escalation mechanisms. If done well, strong third-party governance becomes a competitive advantage in building customer trust.

CIO&Leader: As AI becomes embedded in insurance broking workflows, what governance controls are needed to address risks like model manipulation and data leakage, especially from third-party tools?

MANOJ KERN: AI governance is becoming a critical priority in BFSI, driven by the pace of adoption and the need for controls tailored to AI-specific risks. In insurance broking, where AI powers advisory, comparison, and analytics, the regulatory and reputational stakes are high.

The starting point is a comprehensive inventory of all AI systems, including those within third-party platforms. Many firms lack visibility into vendor AI capabilities, yet accountability remains with them.

To address model manipulation, firms must implement strong validation practices, including adversarial testing, along with continuous monitoring to detect anomalies and bias in outputs.

Data leakage, especially through third-party AI tools, is a major risk. Firms need clear policies on what data can be shared, supported by technical safeguards such as Data Loss Prevention controls to prevent unauthorised exposure.

Vendor governance must ensure transparency around model behaviour, data handling, and compliance obligations, including data residency and explainability requirements.

Equally important is human oversight. AI-driven decisions must remain reviewable and accountable, supported by clear governance structures and, ideally, a dedicated AI risk function.

Ultimately, firms that embed governance into their AI adoption from the outset will be better positioned to innovate responsibly while maintaining regulatory compliance and customer trust.

CIO&Leader: Five years from now, what percentage of insurance advisory and customer interactions do you believe will be AI-assisted, and what does that mean for data security and trust?

MANOJ KERN: To answer this question with the precision it deserves, we must first acknowledge that insurance broking is not a monolithic business. It spans B2B, B2C, and B2B2C models across an extraordinarily diverse range of lines — from Employee Benefits and Cyber to Marine, Marine Hull, Aviation, Satellite, Engineering, Property, Casualty, Liability, POSI, Product Recall, and Reinsurance. The AI adoption trajectory, and its security and trust implications, looks very different across each of these segments. Any answer that treats them uniformly will be both inaccurate and misleading.

With that context, my view is that within five years, the overall percentage of AI-assisted interactions across the broking spectrum will range from 40 to 85 percent — but the distribution will be deeply uneven.

At the higher end of that range sit the B2C and B2B2C segments — retail health, personal lines, SME employee benefits, and group insurance products. These involve high transaction volumes, relatively standardised data inputs, and customer interactions that lend themselves well to AI-driven personalisation, comparison, needs assessment, and

servicing. Here, 80 to 85 percent AI assistance within five years is not only plausible — it is already the direction of travel for the more digitally advanced firms in the market.

At the lower end sit the complex specialty and commercial lines — Aviation, Marine Hull, Satellite, Engineering, POSI, Product Recall, and Reinsurance. These are bespoke, high-value, relationship-driven transactions where the broker's expertise, market knowledge, and negotiation capability are the core value proposition. AI will play an increasingly important role in data aggregation, risk modelling, exposure analysis, and document processing — but the advisory interaction itself will remain predominantly human-led, with AI serving as a powerful analytical layer rather than a customer-facing interface. Here, 40 to 50 percent AI assistance is a more realistic and responsible projection.

In the middle sit lines like Property, Casualty, Cyber, Liability, and standard Employee Benefits — where AI will progressively handle routine interactions, renewal workflows, claims tracking, and regulatory reporting, while complex risk placements and large account advisory work retains a strong human dimension.

Cyber insurance deserves a specific mention because it sits at a uniquely interesting intersection. The very technology that is driving AI adoption in broking — and the threat landscape it creates — is also the subject matter of Cyber insurance itself. AI will be essential in assessing dynamic, rapidly evolving cyber risk profiles. But it also introduces new risk vectors that underwriters and brokers need to understand deeply. In Cyber insurance, AI is simultaneously the tool, the risk, and the subject of the policy.

From a data security standpoint, the diversity of lines means the diversity of data types flowing through AI systems is extraordinary. Employee benefits interactions involve sensitive health and demographic data. Marine and Aviation transactions involve commercially sensitive cargo and asset information. Engineering and Satellite placements involve proprietary technical specifications. Reinsurance transactions involve aggregated portfolio data that could be competitively damaging if exposed. Each of these data categories demands a tailored security posture within the AI architecture — not a one-size-fits-all approach. Firms that deploy AI without calibrating their data security controls to the sensitivity of each line of business are creating asymmetric risk exposure.

"Data leakage, especially through third-party AI tools, is a major risk. Firms need clear policies on what data can be shared, supported by technical safeguards such as Data Loss Prevention controls to prevent unauthorised exposure."

On trust, the stakes vary by segment but are universally high. In B2C segments, customers need to trust that AI recommendations reflect their best interests, not commercial optimisation. In B2B and specialty lines, corporate clients and risk managers need to trust that AI-assisted analysis is accurate, explainable, and not substituting algorithmic convenience for genuine expertise. In Reinsurance, the trust question extends to the integrity and confidentiality of the portfolio-level data that AI systems will inevitably process.

The overarching principle is this: the breadth of insurance broking as a business demands a segmented, risk-calibrated approach to AI governance — not a single policy applied uniformly. The firms that recognise this complexity and architect their AI adoption accordingly will be the ones that earn the trust of their clients, satisfy their regulators, and sustain their competitive advantage across every line they write. ■

Scaling Agentic AI is primarily an operating model challenge

Nitin Mehta, Digital Risk Leader and Consulting Partner, EY India explains why enterprises are shifting from copilots to agentic AI, highlighting governance, security, accountability, and risk frameworks needed to scale autonomous systems responsibly.

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

In an exclusive interaction, Nitin Mehta, Digital Risk Leader and Consulting Partner, EY India, highlights how Indian enterprises are navigating this shift from augmentation to delegation. He examines where agentic adoption is accelerating across functions, the hidden operational costs of managing autonomous systems, and the governance frameworks required to ensure control, auditability, and resilience.

CIO&Leader: How does EY distinguish between the "Copilot era" (generative assistance) and the "Agentic era" (autonomous execution), and what specific triggers are driving Indian enterprises to make this jump now?

NITIN MEHTA: TThe Copilot era is fundamentally about augmentation: AI supports individuals to draft, summarise, code, analyse and retrieve information, while humans still initiate actions and remain accountable for outcomes. The Agentic era represents a move to delegation: AI systems can plan tasks, invoke tools, orchestrate across platforms (for example ITSM, ERP and CRM), and execute multi-step workflows with bounded autonomy—operating within defined policies and producing an auditable record of decisions and actions.

The Agentic era represents a move to delegation: AI systems can plan tasks, invoke tools, orchestrate across platforms and execute multi-step workflows with bounded autonomy.

Indian enterprises are accelerating this shift for a set of converging factors that are increasingly visible at the executive and board agenda:

- sustained cost and service-level pressure in shared services and operations;
- maturity of cloud platforms, APIs and automation foundations (RPA, DevOps, ITSM) that agents can now "plug into";
- sharper competitive cycles in digital channels where speed-to-execution is a differentiator; and
- leadership appetite to move from pilots to measurable outcomes—cycle-time reduction, fewer handoffs, and always-on operations.

In practical terms, the shift is from using AI as an individual productivity enhancer to enabling AI to complete well-bounded activities end-to-end, with human oversight focused on exceptions. This requires explicit decision rights: what the agent may do autonomously, what requires pre-approval, and what must be escalated. .

CIO&Leader: Aside from coding, which enterprise functions are seeing the most aggressive adoption of agentic workflows, and where is the "risk-to-reward" ratio currently most favourable?

NITIN MEHTA: Beyond software engineering, the most rapid adoption is emerging in functions characterised by high transaction volumes, repeatable decision patterns, and strong digital traceability. The underlying differentiator is that agents can combine judgement with execution—progressing from generating content to initiating and completing controlled actions across enterprise systems. Early adoption is most evident in:

- **IT operations & service management** (incident triage, runbook execution, patch orchestration, change validation).

- **Customer operations** (case summarisation, resolution playbooks, next-best-action, proactive outreach with approvals).
- **Finance shared services** (invoice exception handling, reconciliations, collections workflows, close support).
- **Procurement & vendor management** (RFx drafting, compliance checks, contract clause review, supplier queries).
- **Risk & compliance operations** (control testing support, evidence gathering, policy mapping, continuous monitoring signals).

At present, the most favourable risk-to-reward profile is typically found in internal, well-instrumented workflows where data remains controlled, operating boundaries are clear, and actions are reversible—such as ticketing, knowledge-base maintenance, report generation, reconciliation, evidence gathering, and runbook-driven operations. A pragmatic scaling approach is to introduce autonomy progressively: begin with "read + recommend," move to "recommend + execute with approvals," and only then extend to straight-through execution. Conversely, use cases such as autonomous customer communications, pricing decisions, and credit outcomes can carry materially higher conduct, regulatory, and reputational risk unless guardrails, approval pathways and monitoring are already mature.

CIO&Leader: Organizations often cite productivity, but what are the hidden operational costs of managing a fleet of autonomous agents that leaders often overlook?

NITIN MEHTA: While productivity uplift is often the headline benefit, leaders frequently underestimate the operational effort required to run an "agent workforce" at scale. In practice, this resembles establishing a new digital operating capability—supported by controls, monitoring, resilience and continuous improvement—rather than maintaining a conventional software application. Commonly overlooked cost drivers include:

- **Identity, access and credential hygiene:** issuing least-privilege roles, rotation, secrets management, and continuous entitlement reviews.
- **Observability and auditability:** storing agent logs, decisions, tool calls, prompts, and evidence trails—plus analytics to detect abnormal behaviour.
- **Change management for prompts, policies and tools:** versioning "agent instructions," testing in sandboxes, and controlled promotion to production.
- **Data readiness and knowledge curation:** keeping SOPs, runbooks and knowledge bases current so agents don't automate yesterday's process.
- **Model and vendor lifecycle:** monitoring performance drift, cost drift (token/compute), outages, and third-party dependency risk.
- **Human supervision and exception handling:** analysts and SMEs spending time on escalations, approvals, and post-incident reviews.

CIO&Leader: In traditional software, we worry about bugs; in Agentic AI, we worry about "agentic drift"—where an agent takes an unpredictable path to a goal. How do you build a governance framework that monitors intent and pathway rather than just output?

NITIN MEHTA: In agentic systems, governance cannot rely on final-output review alone. What is required is behaviour assurance: the ability to demonstrate that the agent remained within its mandate and followed an acceptable route to the outcome. Leading frameworks therefore monitor three dimensions—what the agent is attempting to achieve, how it is pursuing the objective, and what it ultimately changed in the environment:

- **Intent:** what the agent is authorised to achieve (scope, objectives, prohibited actions) expressed as machine-enforceable policies.

"The Agentic era represents a move to delegation: AI systems can plan tasks, invoke tools, orchestrate across platforms and execute multi-step workflows with bounded autonomy."

- **Pathway:** how the agent is pursuing the goal (plans, tool choices, data sources used, escalation decisions, retries).
- **Impact:** what changed in the real world (transactions, records updated, infrastructure changes, customer communications).

In operational terms, this is typically implemented through controls such as:

- **policy-as-code guardrails** (allowed tools/actions, spend limits, data boundaries);
- **step-up approvals for high-risk actions** (e.g., "write" to production, external communications, financial postings);
- **continuous pathway monitoring** (loops, unusual tool sequences, out-of-hours privileged actions, excessive retries); and
- **traceability**—a defensible record of what the agent observed, the plan it formed, the tools it used, and the outcome. When these are in place, autonomy becomes a controllable spectrum rather than a binary on/off decision.

CIO&Leader: If an autonomous infrastructure agent makes a logic error that leads to a significant system outage, where does the legal and operational liability sit in a post-copilot world?

NITIN MEHTA: In a post-copilot environment, "the agent did it" is not a defensible accountability position. In most operating models, responsibility rests with the enterprise that deployed the agent and defined its permissions—alongside any vendor and service-provider obligations established contractually. For this reason, autonomous agents should be treated as production-grade automation, governed with the same rigour applied to other high-impact change and execution mechanisms.

In a post-copilot environment, 'the agent did it' is not a defensible accountability position.

Practically, leading organisations make accountability unambiguous through:

- an explicit RACI per agent (business owner, technical owner, risk/control owner);
- change-management alignment (approvals, segregation of duties, rollback and post-change validation);
- vendor and model risk controls (SLAs, incident response obligations, audit rights, and clarity on shared responsibility); and
- post-incident forensics (immutable logs of what the agent observed, decided and executed). The goal is not to eliminate error—humans and systems both fail—but to ensure you can detect, contain, explain, and remediate quickly.

CIO&Leader: As agents gain "write" access to databases and infrastructure, they become high-value targets. What are the non-negotiable security guardrails for an agentic system?

NITIN MEHTA: As agents obtain "write" access to enterprise systems, security must prioritise impact of compromise—not only the risk of incorrect outputs. The minimum baseline should mirror privileged human access controls, implemented in ways that are automated, testable and continuously monitored:

- **Least privilege by design:** separate identities per agent, scoped roles, and time-bound elevation for sensitive actions.
- **Strong secrets management:** no hard-coded credentials; rotate keys; isolate tokens; use managed vaults.
- **Tool allowlists and action controls:** agents can only call approved APIs/tools; enforce transaction limits, rate limits, and approval gates for high-impact changes.
- **Network and environment isolation:** sandbox testing, segmented production access, and controlled egress to external sites/services.
- **Prompt and data protection:** prevent prompt injection, restrict untrusted content, and enforce data boundaries (PII, confidential, regulated data).
- **Continuous monitoring:** anomaly detection for unusual actions, excessive retries, and privilege misuse; integrate with SOC workflows.
- **Human-approved "break-glass" and rollback:** safe stop, kill-switch, and automated rollback paths for critical operations.

A useful framing is to treat any powerful agent as a privileged identity that operates continuously. If an always-on administrator account would be unacceptable without strong controls, the same standard should apply to agents: enforce least privilege, make actions fully auditable, and design for rapid containment and rollback from day one. .

CIO&Leader: How should enterprises design "human-in-the-loop" checkpoints without bottlenecking the very speed and autonomy that make agents valuable?

NITIN MEHTA: The objective is to move from 'human reviews everything' to risk-based supervision. Speed is preserved by defining which activities can execute straight-through and which require checkpoints based on impact, uncertainty and control maturity. Effective designs typically anchor oversight around:

- **Impact thresholds:** automatic execution for low-impact actions; approvals only when cost, customer impact, data sensitivity, or production change risk crosses a threshold.
- **Exception handling:** let agents run the "happy path," and escalate only when confidence is low, data is missing, or outcomes deviate from expected ranges.
- **Sampling and post-review:** for moderate risk, execute and then audit a sample with rapid rollback capability.
- **Tiered approvals:** operational

approvals for routine changes; specialist approvals (security/risk/legal) only for defined categories.

When designed well, oversight does not become an approvals bottleneck. Instead, it operates as a structured control function: clear decision rules, automated routing, and human intervention only when warranted. In practice, combining "confidence" and "impact" scoring helps make escalation thresholds predictable, consistent and explainable.

CIO&Leader: Many see risk management as a "handbrake" on innovation. How can a robust risk framework actually accelerate the deployment of Agentic AI by building board-level confidence?

NITIN MEHTA: Risk management is sometimes perceived as a constraint, yet a robust framework can materially accelerate deployment by reducing ambiguity and enabling consistent decision-making. Boards and executive committees tend to slow AI programmes when three questions cannot be answered with confidence: "Is it safe?", "Is it compliant?", and "Who is accountable for the outcome?" A robust framework makes answers to these questions repeatable—so every new use case doesn't restart the debate from scratch.

In practice, the accelerators we see working are:

- **a standardised use-case intake** (data classification, impact assessment, control requirements);
- **pre-approved patterns** (e.g., internal summarisation with no write access; controlled agents with approval gates) that teams can deploy quickly;
- **control libraries** for permissions, monitoring, and human oversight; and
- **transparent reporting**—KPIs for value and KRIs for risk.

The outcome is improved time-to-production because leadership can see autonomy being introduced in a controlled, measurable and auditable manner.

CIO&Leader: Given the evolving global AI regulations (like the EU AI Act), how should Indian enterprises future-proof their agentic deployments against upcoming compliance requirements?

NITIN MEHTA: To future-proof agentic deployments, organisations should align to the direction of travel

"In a post-copilot environment, 'the agent did it' is not a defensible accountability position."

across major regulatory regimes: increased transparency, stronger governance, tighter data protection, and demonstrable human oversight—particularly for higher-risk use cases. For Indian enterprises, a practical approach is to establish a robust baseline early and apply it consistently across business units and geographies, rather than retrofitting controls market by market. Key steps include:

- **Classifying use cases by risk** (customer impact, regulated decisions, safety, critical infrastructure) and applying proportional controls.
- **Building audit-ready documentation:** purpose, data sources, limitations, testing results, and change history for each agent.
- **Ensuring traceability:** logs of agent actions, approvals, and material decisions, with retention aligned to regulatory and business needs.
- **Data governance alignment:** data minimisation, consent and privacy controls, cross-border data handling checks, and secure storage.
- Vendor and model governance: due diligence, contractual controls, and ongoing monitoring of third-party model updates.

For global-facing businesses—especially those serving EU customers or operating through EU entities—the practical aim is to evidence that autonomy is controlled: you can explain what the agent is allowed to do, show how it is monitored, and demonstrate how humans intervene for higher-risk decisions. Even where laws differ by market, expectations are fairly consistent—and building these guardrails early avoids expensive rework later. ■

Infrastructure must be viewed holistically across compute, network, storage

Arun Shetty, CTO & Senior Director, Solutions Engineering at Cisco India & South Asia outlines a shift to AI-ready infrastructure, emphasising purpose-built networks, integrated security, and data platforms like Splunk to help enterprises scale AI from pilots to production with resilience and control.

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

The shift from experimentation to execution is defining the current phase of enterprise AI adoption. As organisations move beyond pilots, the focus is no longer limited to models and algorithms, but on the foundational infrastructure required to operationalise AI at scale. Rising compute demands, data gravity, network complexity, and security risks are forcing CIOs to rethink how IT environments are architected to support always-on, autonomous systems.

In this conversation, Arun Shetty, CTO & Senior Director, Solutions Engineering at Cisco India & South Asia, explains how Cisco is re-engineering AI-ready infrastructure. He highlights the move toward agentic systems, the need for purpose-built networks, and the concept of a "Secure AI Factory," alongside the role of observability, data platforms like Splunk, and integrated security in enabling enterprises to scale AI with resilience, trust, and operational simplicity.

CIO&Leader: Cisco is increasingly shifting toward AI-ready infrastructure that goes beyond hardware. What is the most critical architectural change underway, and how is it enabling Indian enterprises to move from AI pilots to production?

ARUN SHETTY: If you look at how AI is evolving, it has moved beyond simply answering questions. We are now transitioning into a phase where AI systems can take actions. This marks the emergence of AI agents, and in the future, we will likely see the rise of physical AI as well. This shift is already underway and is fundamentally changing how enterprises operate.

There are three major changes taking place, along with corresponding constraints.

From a technology perspective, there will be a significant increase in traffic, a rise in the number of devices, and an expansion of risk exposure. AI agents will effectively become part of the workforce, contributing to this growth in scale. As traffic and device density increase, the associated risks also grow proportionally.

From an operational standpoint, complexity is becoming a critical challenge. Even today, enterprise environments are highly complex, with a mix of cloud infrastructure, legacy applications, and varied infrastructure strategies across organisations. AI further amplifies this complexity. As complexity increases, it also exposes a gap in skills, which becomes an operational challenge for enterprises.

The third dimension is people. Expectations from both employees and customers are rising significantly. There is an increasing demand for immediate outcomes and seamless experiences, driven by what AI is now capable of delivering. This shift in expectations is a major factor enterprises must prepare for.

These three dimensions—technology scale, operational complexity, and rising expectations—define the broader impact of AI and how organisations must respond.

"We are now transitioning into a phase where AI systems can take actions. This marks the emergence of AI agents, and in the future, we will likely see the rise of physical AI as well."

In addition to these changes, there are three major constraints or architectural considerations that enterprises must address in the AI era.

The first constraint is infrastructure. This includes power, compute, and networking. These foundational elements must scale to support AI workloads, making infrastructure readiness a primary requirement.

The second constraint is trust, which encompasses both security and safety. Security concerns are widely understood, but safety is equally critical and often less discussed. Safety relates to the intrinsic behavior of AI models and applications. When an input is provided to an AI model, the output is not always consistent. Issues such as hallucinations, toxicity, and unpredictable responses can occur. These behaviors are inherent to how models function, which makes it essential to ensure that models behave in alignment with intended outcomes. This is why safety becomes a critical component alongside security, contributing to what can be described as a trust deficit.

The third constraint is data. Most AI models today are trained on publicly available data, including text, video, and audio. However, enterprises possess their own proprietary data, which can be leveraged to derive more meaningful outcomes. By using enterprise data, organisations can train, distill, and correlate models more effectively, thereby unlocking greater value from AI implementations.

Taken together, these infrastructure, trust, and data challenges represent the critical areas that enterprises must address as they move from AI experimentation to production at scale.

CIO&Leader: Organisations are heavily investing in AI infrastructure, often prioritising compute at scale. From a systems engineering perspective, how should CIOs balance power and data demands for AI while integrating with existing legacy infrastructure?

ARUN SHETTY: From an infrastructure standpoint, the starting point is always the use case. Enterprises must clearly define what they aim to achieve. Based on that, they need to decide whether workloads should be deployed on-premises or in the cloud. In many cases, organisations initially experiment in the cloud and later bring workloads back on-premises due to concerns around data security and data sovereignty.

Identifying the use cases and aligning the required infrastructure to support them is therefore a critical first step. Once this clarity is established, enterprises can determine the number of GPUs required, the level of compute needed, and the overall architecture.

At Cisco, this approach is defined as a Secure AI Infrastructure. This is not limited to compute alone; it is a full-stack architecture developed in collaboration with NVIDIA. The stack ensures end-to-end observability and embeds security across every layer—from silicon to applications. Security is not an add-on but an integral part of the entire architecture.

These architectures are validated through two frameworks. NVIDIA provides certification through its Enterprise Reference Architecture, while Cisco offers its own Cisco Validated Designs (CVDs) to ensure that the infrastructure is optimised and deployment-ready.

Once the infrastructure is established for pilot workloads, the next step is scalability.

The first stage is scale-up, where the existing stack is expanded to increase compute capacity.

The second stage is scale-out, where infrastructure extends beyond a single rack to multiple racks within a data center.

The third stage is scale-across, which becomes necessary when local constraints such as power availability limit further expansion. In such scenarios, workloads are distributed across multiple data centers, often located closer to power sources.

Infrastructure must also be viewed holistically. It includes not

only compute, but also network and storage. Cisco works with ecosystem partners to deliver integrated storage solutions, while leveraging the broader NVIDIA architecture to ensure successful deployment of AI use cases.

On the networking side, high throughput is critical. Cisco's Silicon One G300 enables switching capacity of up to 102.4 Tbps within data centers, while the P200 routing platform supports up to 51.2 Tbps for inter-data center connectivity. These capabilities are essential because AI workloads—particularly inferencing driven by agents—operate continuously and require sustained high bandwidth. As organisations scale across data centers, high-speed interconnects become equally important.

In addition, Splunk plays a key role through its Data Fabric. It enables the ingestion and correlation of large volumes of data, which can then be leveraged for advanced analytics, including MachineGPT, to unlock additional use cases and insights.

This integrated approach defines what Cisco refers to as a Secure AI Factory, where infrastructure is designed with both performance and security as foundational principles.

From a security and safety perspective, protection must extend beyond the infrastructure to the AI applications themselves. Cisco AI Defense addresses this requirement by enabling organisations to discover all AI applications in use, detect vulnerabilities within them, and assess risks associated with downloaded models.

It also provides runtime protection, ensuring that AI applications remain secure while operating. This end-to-end capability—from discovery to runtime protection—forms a critical component of the Cisco Secure AI Factory.

"Organisations must treat agents as digital employees, with defined identities, access controls, and traceability."

CIO&Leader: With the integration of Splunk, how is Cisco re-engineering its networks to function as primary sensors for AI-driven threats?

ARUN SHETTY: To understand this, it is important to first examine the role Splunk plays in completing the broader Cisco solution.

At the core is the concept of digital resilience, which is the ability of an organisation to remain securely operational despite disruptions. These disruptions could range from IT outages and security breaches to unforeseen incidents such as configuration errors. In some cases, incidents cannot be prevented, making the ability to respond effectively just as critical as prevention. As a result, digital resilience has become a board-level priority for organisations.

Enterprise environments today are inherently complex. They span private cloud, public cloud, legacy applications, and modern microservices. Additionally, organisations often rely on infrastructure they do not fully control, such as the internet or SaaS platforms. This lack of ownership leads to limited end-to-end visibility, which in turn delays issue detection and resolution.

To address this, Cisco focuses on three key pillars to achieve digital resilience.

The first is assurance, which relates to end-to-end connectivity across the entire digital ecosystem—from on-premises environments to the cloud. Systems continuously generate telemetry data, which can be collected, distilled, and correlated to identify the exact source of an issue. With complete visibility, organisations can determine whether a problem originates within their own infrastructure or from an external service.

The second pillar is observability, which ensures a consistent and high-quality user experience. By maintaining end-to-end visibility, organisations can monitor application performance, reduce downtime, and proactively address issues before they impact users.

The first constraint is infrastructure. This includes power, compute, and networking."

The third pillar is security operations, which involves the ability to prevent, detect, investigate, and respond to threats. The shift here is from a prevention-only mindset to a comprehensive response-driven approach.

This is where Splunk becomes critical. It acts as a unified data platform, aggregating telemetry from across the enterprise. Different teams—IT, security, and engineering—may use different tools, but they operate on the same underlying data. This shared data foundation simplifies problem identification and accelerates resolution.

The integration of Splunk into Cisco's architecture enables all data generated from Cisco platforms to be ingested and analysed within a single framework. This is referred to as the Cisco Data Fabric, which allows organisations to derive actionable insights and make informed decisions in real time.

In the context of AI, particularly with the rise of AI agents, the requirement shifts toward detecting and responding at machine speed. Traditional response times are no longer sufficient. Organisations must be able to process signals, identify threats, and take action in real time.

This is driving the evolution toward Agentic Security Operations. One aspect of this is enabling advanced analytics and design capabilities that allow organisations to move from reactive to proactive security models. Another is the expansion of the Security Operations Center (SOC) through AI agents.

These agents can take on roles such as detection assistants, detection agents, response orchestrators, and autonomous response agents. They automate workflows, assist human operators, and in some cases, take independent actions based on predefined policies.

This shift represents the next phase of security operations—where AI agents augment human capabilities, improve efficiency, and enable faster, more intelligent responses. It reinforces the role of Splunk not only in digital resilience but also as a foundational platform for AI-driven security in modern enterprise environments.

CIO&Leader: Many AI initiatives fail to scale due to data center limitations rather than model performance. What are the key blind spots that prevent AI projects from scaling?

ARUN SHETTY: This is a critical point. The network plays a central role in AI environments, particularly with the rise of AI agents. Unlike traditional workloads, agents operate continuously, which creates a persistent and consistent load on the network. As a result, infrastructure must be designed to support 24/7 operations.

The first requirement is a purpose-built network for AI. This includes high-speed Ethernet and advanced switching capabilities. For example, emerging architectures can support speeds ranging from 800 Gbps to 1.6 Tbps per port within the LAN environment. Such capabilities are essential to handle the scale and performance requirements of AI workloads.

The second requirement is operational simplicity. Enterprises must have the ability to quickly identify where issues occur and resolve them efficiently. Without clear visibility and simplified operations, scaling becomes difficult.

The third requirement is integrated security. Security must be embedded directly into the network fabric. Given the scale of AI-driven traffic, the network itself must be inherently secure and capable of scaling without introducing vulnerabilities.

At Cisco, these challenges are addressed through the Silicon One architecture and integration with NVIDIA Spectrum-X. This ensures that infrastructure limitations in AI environments are mitigated through optimised performance, scalability, and speed.

CIO&Leader: As AI agents become integral to enterprise environments, how does a self-healing, intelligent network operate under the load of autonomous systems?

ARUN SHETTY: One of the primary challenges in enterprise environments today is the high mean time to detect (MTTD) and mean time to repair (MTTR), largely due to limited visibility. Improving visibility is therefore essential.

The industry is moving beyond traditional monitoring and centralised management toward Agentic Operations (Agentic Ops). This represents a shift from passive monitoring to active execution, where AI agents play a key role.

At Cisco, we introduced AI Canvas as part of this approach. AI Canvas enables three core capabilities.

First, multi-domain troubleshooting. When an issue arises, it may originate from the network, application, server, or other infrastructure components. A self-healing system must be capable of diagnosing issues across all these domains.

Second, collaborative operations. In many enterprises, troubleshooting is performed in silos across different teams. By providing all teams with access to a unified data view, collaboration improves and problem resolution becomes more efficient.

Third, the use of a proprietary Deep Network Model (DNM). This model enables AI-driven diagnostics and automation. For example, an operator can initiate a troubleshooting request, after which AI agents analyse network performance, identify the root cause, and recommend corrective actions.

In practice, the system can go further. With appropriate governance, agents can autonomously execute corrective actions, such as reconfiguring network parameters to resolve performance issues. The human-in-the-loop model remains available, but organisations can also enable fully autonomous operations where appropriate.

These capabilities extend across domains. If an issue is not network-related but originates in the application layer, the system can identify and address it accordingly.

Agentic Ops represents a shift toward simplified operations in highly complex environments, enabling organisations to move from reactive troubleshooting to proactive and autonomous resolution.

CIO&Leader: Beyond encryption, what safeguards must be embedded at the network layer to safely deploy autonomous AI systems?

ARUN SHETTY: The adoption of AI agents significantly expands the attack surface. Addressing this requires a structured approach across three dimensions.

The first is protecting the enterprise from AI agents. Since agents can operate autonomously, it is essential to ensure that they act within defined boundaries. This requires extending Zero Trust principles to AI agents. Organisations must treat agents as digital employees, with defined identities, access controls, and traceability. Each agent should have a human owner, and all actions must be auditable. Strict identity and access management ensures that agents only perform authorised actions.

Organisations must treat agents as digital employees, with defined identities, access controls, and traceability.

The second is protecting AI systems from external threats. This includes both model-level risks and external manipulation attempts. Cisco AI Defense addresses this by enabling discovery of all AI applications, conducting vulnerability assessments, and validating models before deployment. It also includes supply chain security checks for downloaded models and runtime guardrails to ensure safe operation.

The third is detecting and responding at machine speed. As discussed earlier, platforms like Splunk enable real-time detection and response through Agentic Security Operations. This ensures that threats are identified and mitigated at the speed required in AI-driven environments.

Together, these layers, identity, protection, and rapid response, form the foundation of secure AI deployment.

CIO&Leader: Compared to the cloud transition, what makes the AI infrastructure shift more complex for modern CIOs?

ARUN SHETTY: The defining factor is the pace of innovation. AI is evolving at a significantly faster rate than previous technology shifts.

To illustrate, the United States is projected to require approximately 62 GW of power for AI workloads by 2028. Context window sizes are expanding rapidly, currently reaching around one million tokens. At the same time, global AI spending is expected to reach approximately US $3.3 trillion.

AI has the potential to either disrupt or accelerate every industry. Organisations must therefore proactively identify relevant use cases and align their strategies accordingly.

From an enterprise perspective, this requires a structured approach to identifying opportunities, deploying use cases, and scaling infrastructure. At the same time, security and safety remain critical priorities, requiring continuous visibility and governance.

In the Indian context, the scale of transformation is equally significant. Data center capacity is expected to reach approximately 8 GW by 2030, with infrastructure investments projected to reach US $70 billion by 2026. AI is also expected to contribute around US $1.7 trillion to the Indian economy.

These indicators highlight the magnitude of the shift underway. AI will fundamentally reshape how organizations operate, both by introducing disruption and by accelerating existing processes. Enterprises must be prepared to navigate both outcomes. ■

If your data isn't real-time, your business is already behind

Andrew Sellers, VP – Technology Strategy and Enablement, and Rubal Sahni, AVP – India and Emerging Markets, Confluent, discuss why data streaming is becoming a foundational layer for real-time agility, AI, and enterprise transformation.

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

Confluent, the California-based data streaming pioneer built around Apache Kafka, is at the forefront of helping enterprises rethink how they harness real-time data. Originally developed as a distributed event streaming backbone, Kafka has evolved from a messaging system into a foundational layer for modern, event-driven architectures—enabling organizations to continuously capture, move, and process data as it is generated.

At a time when CIOs are under pressure to scale AI beyond pilots, deliver real-time experiences, and prove tangible ROI, the need for reliable, governed, and continuously available data has become critical. Platforms like Confluent are extending Kafka's capabilities into a full-fledged data streaming ecosystem, allowing enterprises to move beyond batch processing toward always-on, intelligent operations.

On the sidelines of the Confluent Data Streaming World Tour in Mumbai, Andrew Sellers, VP – Technology Strategy and Enablement, and Rubal Sahni, AVP – India and Emerging Markets, spoke with CIO&Leader about why data streaming is no longer just about pipelines, but a foundational layer for agility, AI, and real-time enterprise transformation.

CIO&Leader: Data streaming is often seen as just pipelines. What's changing in how enterprises should think about it?
ANDREW SELLERS: That view is rapidly evolving. Data streaming is no longer just about moving data in real time; it's about re-architecting how enterprises operate.

Having served as a CTO twice, what drew me to streaming was its ability to decouple teams, systems, and technologies. Beyond real-time movement, it enables continuous data reuse, accelerating time to market for new applications. Once data is produced, it can be leveraged repeatedly, driving faster innovation. Digital-native firms, especially in India, build this in from the start, while legacy enterprises use it to break silos and modernize data into lakes.

For CIOs and CTOs, the real value isn't cost efficiency—it's speed and agility. When data is discoverable, contextualized, and trusted, building new applications becomes significantly faster. In most modern use cases, the bottleneck isn't model development but accessing reliable data inputs. Streaming addresses this foundational challenge, enabling innovation at scale.

CIO&Leader: How does this play out differently for startups versus large enterprises?
ANDREW SELLERS: The contrast is quite stark. Digital-native startups have the advantage of starting from a clean slate. They are not constrained by legacy systems, so they can adopt modern architectures—including streaming—from day one. This allows them to move faster and innovate without friction.

"For CIOs, the true value of streaming isn't cost efficiency, it's speed, agility, and the ability to innovate continuously."

Large enterprises, however, typically begin with targeted use cases. These are often centered around:

- Breaking data silos
- Enabling real-time decision-making
- Feeding operational data into analytics platforms

Once they see value in a single use case, something important happens—they develop confidence and operational muscle memory. From there, adoption expands organically across the enterprise.

The most successful transformations don't start with large, risky bets. They start small, prove value, and then scale. We don't push large upfront costs. Instead, we let them fall in love with the first use case. Once they see success, they come back for more. Our biggest accounts all started with one use case. Executive buyers really appreciate that approach because they stake their careers on these investments, so small steps build trust.

CIO&Leader: We are seeing industries move toward instant experiences, quick commerce, real-time lending, AI-driven decisions. How critical is real-time data in this shift?

RUBAL SAHNI: It's absolutely foundational.

If you look at the evolution of technology in business:

- A decade ago, it was an enabler
- Five years ago, it became a competitive advantage
- Today, it is a necessity for survival

Industries are being reshaped by speed of decision-making. Whether it's delivering products in hours or approving loans in minutes, the expectation is shifting toward immediacy.

This is only possible when systems can: Process data in real time; Make intelligent decisions instantly and continuously adapt based on new information

Streaming acts as the backbone for this shift. Without it, achieving true real-time intelligence at scale is extremely difficult.

CIO&Leader: Despite all the progress, many organizations struggle to move AI from pilot to production. What's the real bottleneck?

ANDREW SELLERS: It's not the models—it's the data infrastructure. Most organizations already have capable models; the challenge is ensuring the data feeding them is accurate, contextual, secure, and continuously available.

Production AI also demands confidence and control. Enterprises need to understand why a decision was made, whether it can be reproduced, and how it can be audited. Yet many modern AI frameworks prioritize speed over transparency, with a single request triggering multiple opaque processes—making behavior hard to track and costs difficult to manage.

What's needed is a system that

captures the full flow of data and decisions, enabling visibility, traceability, and continuous improvement. Without that, scaling AI remains risky. Building a trusted data environment—where data is observable, auditable, and aligned—is essential to giving AI applications a reliable foundation.

CIO&Leader: Can you illustrate the risks of weak data infrastructure with a real example?

RUBAL SAHNI: Let me answer this. A recent case highlights this clearly. An OCR system misread Rs 80,000 as Rs 80 lakh due to a formatting issue. The payment was processed instantly, and the error was only discovered days later.

The consequences went beyond financial loss. There were:

- Compliance and taxation complications
- Reputational damage
- Customer trust issues

The bigger problem was the lack of visibility—there was no easy way to trace how the error occurred or prevent it from happening again.

This is where real-time data infrastructure with governance becomes critical. It enables organizations to detect anomalies instantly, trace root causes, and enforce safeguards before issues escalate.

CIO&Leader: Data streaming is often seen as just pipelines. What's changing in how enterprises should think about it?

ANDREW SELLERS: That view is rapidly evolving. Data streaming is no longer just about moving data in real time; it's about re-architecting how enterprises operate.

CIO&Leader: What are the most common mistakes enterprises make when scaling data streaming platforms like Kafka?

ANDREW SELLERS: One of the biggest misconceptions is that open-source solutions are inherently "free." They come with significant operational complexity.

Enterprises often struggle with capacity planning. They are either:

- Over-provision for peak workloads, leading to high costs
- Under-provision, resulting in outages and performance issues

Another challenge lies in architectural decisions, such as partitioning, which can be difficult to reverse once implemented. What organizations need is an approach that abstracts this complexity—allowing them to scale dynamically, optimize costs, and maintain reliability without deep operational overhead. More broadly, the shift is from treating streaming as a tool to treating it as a platform, with integrated capabilities around processing, connectivity, and governance.

CIO&Leader: What partitioning issues do startups and large enterprises face differently with Kafka?

ANDREW SELLERS: Partitioning decisions are hard. You must define partitions early, but adding partitions later breaks ordering guarantees. Startups and large enterprises both struggle: startups often don't anticipate scaling needs, while large enterprises must plan for peak workloads. Our tools help them balance this—ensuring flexibility, scaling as needed, and maintaining high availability.

CIO&Leader: Why is data governance a blind spot for many enterprises scaling AI?

ANDREW SELLERS: Data governance boils down to metadata creation, but developers don't like writing metadata, they prefer to move fast. Yet, governance depends on metadata. AI can help by automating schema induction, it looks at samples and suggests metadata, which developers can quickly validate. This is critical because AI governance depends on strong data governance. If you don't govern data, you can't govern AI.

CIO&Leader: Finally, what's your view on Agentic AI—are we still early, or is it already taking shape?

ANDREW SELLERS AND RUBAL SAHNI: It's already taking shape, and faster than many expected.

Six to eight months ago, most enterprises were experimenting with proofs of concept. Today, a growing number are running agentic systems in production.

Agentic AI exists on a spectrum—from structured automation to more autonomous, decision-making systems. What's changing is that these systems are now being deployed in real business workflows.

However, their success depends heavily on the underlying data infrastructure. Without real-time, governed, and traceable data, agentic systems can quickly become unpredictable. So, while the momentum is real, the organizations that will succeed are the ones that invest not just in AI models, but in the data foundations that make AI reliable at scale. ■

NxtGen

presents

ET Edge

CISO FORUM CONFERENCE '26 MID YEAR

Building A Cyber Resilient Enterprise

12 -14 June 2026 ▪ Hayatt Regency Jaipur Mansarovar

3-Day Residential Experience in the regal legacy of Jaipur with 100+ India's top CISOs.

Key Conference highlights

- 3 Day Residential Experience
- 11 Hours of structured face-time with CISOs
- 1:1 Meetings with CISOs buying cycle aligns with your solution
- Cultural evening with Rajasthani Folk group

Eminent Speakers

Krishnamachari Srikkanth
Former Indian Cricketer

Kiran Gopinath
Chief Innovation Officer & Head, Sahamati Labs

Shashikant Dahuja
Executive Director and Chief Underwriting Officer, Shriram General Insurance

Munish Chandan
Addl. CITO & Dy. CISO, Citizen Resources Information Department, Govt. of Haryana

For Sponsorship, Write to:

Hafeez Shaikh
Assistant Director - Projects, ET Edge
hafeez.shaikh@timesgroup.com,
+91 9833103611

Supriya Sahoo
Senior Project Manager,
supriya.sahoo@timesgroup.com,
+91 8095056886

PRESENTING PARTNER: NxtGen | GOLD PARTNER: ARMIS from ServiceNow | SILVER PARTNERS: airtel, netskope, proofpoint, CYBERPWN, rubrik | STRATEGIC CASE STUDY PARTNER: SECLORE

ASSOCIATE PARTNER: orca SECURITY | EXHIBIT PARTNERS: Barracuda, CyberAssure, cybermindr, HarkX, MATTERS, CyberAssure, PrivacyPillar, techsa | CONCEPT BY: ET Edge CISOFORUM

RNI No. DELENG/2012/42736

AMD
presents

ET Edge 27th Annual Conference

CIO&LEADER

The Agentic Enterprise

PLATFORMS. PEOPLE. POLICY. PROFITS.

co-presented by NxtGen

#CIOandLeaderConference

30 JULY – 02 AUGUST 2026

FAIRMONT, JAIPUR

A 4-Day Experiential Retreat with 200+ of India's Top CIOs & Future CIOs.

WHAT's IN FOR YOU

▶ 1:1 meeting with CIOs ▶ Thought Leadership stage time with 200+ CIOs taking notes
▶ Ideas Cafe, Roundtable, Workshop, Booths ▶ 2160+ Networking minutes
▶ CIO Samman & NEXT100 Awards ▶ Evening with leading sports & bollywood celebrity

EMINENT SPEAKERS

Shri Rana Ashutosh Kumar Singh
MD (International Banking, Global Markets & Technology), State Bank of India

Madan Sunder Das
Monk, ISKON, Founder, Spiritual Guide and Community Leader, EVOLVE Pune

Be Where India's Tech Leaders Meet.

PRESENTING PARTNER AMD | CO-PRESENTING PARTNER NxtGen | GOLD PARTNER dynatrace | ASSOCIATE PARTNERS DELTA, TECHNO DIGITAL | EXHIBIT PARTNER BARCO ClickShare | CONCEPT BY ET Edge CIO&LEADER, NEXT100

For Sponsorship, Write to

Hafeez Shaikh
Assistant Director - Projects
hafeez.shaikh@timesgroup.com
+91 9833103611

Supriya Sahoo
Senior Project Manager
supriya.sahoo@timesgroup.com
+91 8095056886

#CIOandLeaderConference