CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

APRIL 2026 • VOLUME 15 • ISSUE 01 • ₹150

CIO PLAYBOOK 2026

SET CONTEXT. SHAPE DECISION. WIN BIG.

India's top 100 CIOs convened in Varanasi for a 3-day residential forum, defining the technology agenda across AI, data, security, and operations.

Anil Shankar, Global CIO, Tata International; **Ashish Kshetry,** CIO & VP – IT, Asian Paints; **Chitti Babu,** Group CIO, Aurobindo Pharma; **Dipu K.V.,** Senior President, Bajaj General Insurance; **Gaurav Duggal,** CIO, Jio Platforms; **Lalita Chandel,** Global Consulting Practice – CIO Advisory, TCS; **Mayank Sharma,** Chief AI & Innovation Officer, IIFL Finance; **Nikhil Bhushan,** CTO, Tata Starbucks; **Ninad Raje,** Group CIO, Times Network; **Narrottam Sharma,** CIO, Jubilant FoodWorks; **Prakash Dharmani,** CIO, Okaya Power Group; **Pratik Pal,** Advisor – Data and AI, Chairman's Office, Tata Group; **Ramesh Narayanswamy,** Group President – Digital and Data Intelligence, Hinduja Group; **Ranjitha R.,** Head of Engineering, PhonePe, among others.

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

The essential ingredients of AI success

Indian enterprises are rapidly accelerating their efforts to experiment with and deploy AI at scale. Over the past few years, there has been a surge of initiatives aimed at staying ahead of the curve. Yet, as organizations move from experimentation to operationalization, the real challenges are beginning to surface—stretched infrastructure, performance bottlenecks, unclear ROI, and growing resistance from both customers and employees.

According to Gartner, by the end of 2025, at least 50% of generative AI projects were abandoned after the proof-of-concept stage, undone by poor data quality, weak risk controls, escalating costs, and a lack of tangible business value.

This is the inflection point. As enterprises double down on agentic AI, success will depend not on isolated experimentation, but on a deliberate, long-term playbook. That means investing in strong data and storage foundations, enabling enterprise-wide change management, and embedding robust guardrails around safety, privacy, accountability, and fairness. Just as importantly, organizations must define, and rigorously track clear, measurable business outcomes.

As Ninad Raje, Group CIO at The Times Group, aptly observed at the CIO Playbook conference: "If you have a Ferrari but are driving it on Mumbai's pothole-ridden roads, it won't work, it will crack open." The message is clear: AI failures are rarely about the models. More often, they stem from gaps in data, infrastructure, and organizational readiness.

Notably, startups and digital-native companies seem to be navigating this shift more effectively leveraging cloud-native architectures, agile ways of working, and a sharper focus on outcomes from the outset.

In this edition, we bring you highlights from the CIO Playbook conference held in March, where the country's top CIOs shared candid perspectives on what it truly takes to move AI from promise to performance, and to deliver sustained business impact. I look forward to hearing your perspective on what it truly takes to deploy AI at scale. ■

"As enterprises double down on agentic AI, success will depend not on isolated experimentation, but on a deliberate, long-term playbook."

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

CONTENT APRIL 2026

COVER STORY

10-28

The CIO Playbook 2026

The CIO Playbook 2026, held in Varanasi, brought together India's top 100 CIOs to redefine enterprise strategy and advance AI, data, and security agendas for the next decade of transformation.

Cover Design by:
Shokeen Saifi

Please Recycle this Magazine and Remove Inserts before Recycling

NEWS & VIEWS

INSIGHT

TECH TALK

CIO&LEADER

www.cioandleader.com

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

EDITORIAL
Group Editor: **R Giridhar**
Editor: **Jatinder Singh**
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

Ramesh Narayanaswamy Appointed Group President – Digital & Data Intelligence at Hinduja Group

Ramesh Narayanaswamy has been appointed to lead enterprise digital initiatives, data intelligence capabilities, and technology-driven transformation across the Hinduja Group's diverse global business portfolio.

Aman Malhotra Appointed Security Solution Engineering Leader at Microsoft

Aman Malhotra has taken on the role at Microsoft, where he will lead security solution engineering across India and South Asia, strengthening cyber resilience and accelerating AI-driven security adoption.

Cijo George Appointed VP & Head of AI at Practo

Cijo George has been appointed Vice President & Head of AI at Practo, where he will lead the company's artificial intelligence strategy, platform development, and machine learning–driven healthcare innovation initiatives.

Jeetendra Dhakane Appointed Head – IT & Digital Transformation at RP Tech India

Jeetendra Dhakane has taken charge as Head – IT & Digital Transformation at RP Tech India, where he will lead enterprise IT strategy, digital modernization, and technology-driven transformation initiatives across the company's operations.

NEXT100 Winner Alok Kumar Appointed GM – IT at Shahi Exports

Alok Kumar has joined Shahi Exports as General Manager – Information Technology, where he will oversee enterprise IT strategy, infrastructure, and digital transformation initiatives.

Suman Guha Appointed Chief Digital Officer at Croma

Suman Guha has joined Croma as Chief Digital Officer, where he will lead digital transformation, omnichannel strategy, and the development of a scalable, future-ready technology architecture for the retail business.

Reema Jain Appointed Chief Information Officer at Unilever

Reema Jain has been appointed Chief Information Officer at Unilever, where she will lead global technology strategy, enterprise architecture, and digital transformation across markets.

Arjun Laxminarasu Appointed Head – Marketing at UiPath

Arjun Laxminarasu joins UiPath as Head of Marketing – India and South East Asia to drive demand generation, brand positioning, and regional growth initiatives.

Arvind Prithvinath Singh Appointed CIO at M3M India

Arvind Prithvinath Singh has been appointed Chief Information Officer at M3M India, where he will lead enterprise IT strategy, digital innovation, and technology-driven transformation initiatives.

Abhoy Kumar Sarkar Appointed MD & Country Manager at Redis

Abhoy Kumar Sarkar has been appointed Managing Director and Country Manager for India and SAARC, where he will drive regional growth and an AI-led data innovation strategy.

Shashwat Sharma Appointed MD & CEO at Airtel India

Shashwat Sharma has been appointed Managing Director and Chief Executive Officer at Airtel India, where he will lead operations, growth strategy, and customer-centric innovation initiatives.

Bhawesh Thakar Appointed CDIO at Wagh Bakri Tea Group

Bhawesh Thakar steps into the role of Chief Digital & Information Officer to lead digital transformation, IT strategy, and enterprise technology modernization initiatives.

Vinay Jain Appointed Chief Information Security Officer at Kiwi Insurance

Vinay Jain has been appointed Chief Information Security Officer at Kiwi General Insurance, where he will lead cybersecurity strategy, risk governance, and secure digital transformation initiatives.

Aditya Singh Elevated as Chief Customer Experience Officer at ACT Fibernet

Aditya Singh has been elevated to Chief Customer Experience Officer at ACT Fibernet, where he will lead customer experience transformation, service delivery, and customer retention initiatives across the company's operations.

Himanshu Bhardwaj Appointed Chief Information Officer at Adani Healthcare

Himanshu Bhardwaj has been appointed Chief Information Officer at Adani Healthcare, where he will lead enterprise IT strategy, digital health initiatives, and technology-driven transformation across operations.

Why AI means growth for Wipro and cuts for Oracle?

Wipro and Oracle take divergent AI paths: Wipro hires 50,000 to scale AI services and talent, while Oracle cuts 12,000 jobs to fund data centres and capital-intensive cloud infrastructure.

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

Wipro and Oracle have made into headlines with contrasting workforce decisions. Wipro announced plans to hire 50,000 new employees to expand its service capabilities. Meanwhile, Oracle moved to reduce its headcount by approximately 12,000 roles in India alone.

While both companies cite artificial intelligence as their primary driver, their methods reveal a fundamental difference between a service provider and a product manufacturer.

Building the human infrastructure for AI services

Wipro's decision to add 50,000 roles aligns with its 'ai360 ecosystem'. This US $1billion initiative focuses on integrating AI into every tool and client solution.

AI is splitting the tech industry in two: services firms scale people, while product giants like Oracle cut jobs to fund machines.

The company is focusing on agentic AI systems that can take independent action to solve business problems. They are currently upskilling their entire 250,000- person workforce while hiring fresh talent to meet the rising demand for custom AI builds in sector like healthcare and finance.

Shifting capital to hardware and automation

Oracle's approach is different because its core business involves building the physical and digital infrastructure that AI runs on. Oracle is racing to build massive data centres to compete with giants like Amazon and Microsoft.

To fund this expansion Oracle is streamlining its operations. The company identified 12,000 roles in India as 'redundant', many of which were in traditional sales, marketing and back-office functions. By reducing its payroll, Oracle frees up cash flow to pay for high-end Nvidia chips and power-intensive data facilities.

Key differences in strategy

The divergence comes down to where each company sits in the technology stack. Wipro acts as the "architect." They help businesses figure out how to use AI. Because every client has different needs, Wipro requires a large, skilled workforce to provide personalised service. Hiring 50,000 people is a bet that the market for AI advice and implementation will continue to grow.

While Oracle acts as the "utility provider." They provide the electricity (cloud computing and databases) that powers AI. This is a high-cost, high-margin business that favours automation over large headcounts. Oracle is willing to cut 12,000 jobs today to ensure it has the billions of dollars needed to dominate the cloud market of 2030. ■

India is betting US $1.3 billion on AI

A sweeping new Parliamentary report lays bare India's ambitious technology gamble from battlefield drones to deepfake laws, and 47 lakh new tech jobs.

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

India's Parliament doesn't often make headlines in Silicon Valley. But the Twenty-Seventh Report, presented on March 30, 2026, should land on every tech investor's desk. Covering AI's impact, it reads less like a committee report and more like a national battle plan, detailed, urgent, and candid about gaps.

The numbers that matter

The report opens with an audacious forecast: AI could add $967 billion to India's economy by 2035 and up to $500 billion to GDP by 2025, accounting for 10% of the $5 trillion target. Adoption has risen from 8% in 2023 to 25% in 2024. At the center is the IndiaAI Mission, approved in March 2024, with a budget of Rs. 10,371 crore over five years across seven pillars, including compute, datasets, innovation, skilling, startups, and safe AI.

Building India's own AI brain

One of the report's striking revelations is India's push for AI sovereignty. Twelve startups, including Sarvam AI, Soket AI, Gnani AI, and BharatGen, are building LLMs on Indian data and servers. Over 38,000 GPUs are subsidized under Rs. 65 per hour. A 3,000-chip cluster is under construction, though dependence on imported processors remains.

AI on the battlefield and the border

Perhaps the most consequential section concerns defense. DRDO is integrating AI into surveillance, missile guidance, autonomous drones, and cyber warfare. Projects include face recognition under disguise, Mandarin-to-Hindi translation, sonar classification, and AI fencing. India's forces are developing "RMA 3.0." The committee flags fragmented data and costly retrofitting, urging a Defense Data and Compute Backbone.

The deepfake problem gets a legal answer

On the civilian side, the report marks a regulatory milestone. In February 2026, India notified IT Amendment Rules targeting Synthetically Generated Information (SGI). Platforms must label AI content, verify users, and respond to takedowns within 2 hours for serious harms. The committee raised concerns about deepfakes targeting women.

The jobs question — answered honestly

The report doesn't shy away from workforce anxiety. It cites estimates suggesting that 38 million jobs could be displaced, while AI could create 47 lakh new tech roles by 2027. Quoting Prime Minister Modi: "Loss of jobs is AI's most feared disruption, but history shows work does not disappear, only its nature changes." To bridge the gap, India is setting up 570 AI Data Labs, training 1.68 lakh individuals, and partnering with Microsoft to train 5 lakh people by 2026.

The bottom line

This Parliamentary report is a serious, candid document. It acknowledges India's strengths, the world's largest AI skill penetration, 900 million internet users, and a booming startup ecosystem, while being frank about gaps in data, compute, and defense readiness. ■

Wipro signs US $1 Bn Olam deal, acquires Mindsprint

Wipro's $1B Olam Pact and $375M Mindsprint Acquisition Mark a Strategic Push into AI-Led Agri-Tech, Securing Long-Term Farm-to-Fork Transformation Deals and Deepening Global Supply Chain Expertise

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

Wipro signed a definitive agreement on April 5, 2026, to acquire Mindsprint, the IT services arm of Singapore-based Olam Group. The cash transaction carries a purchase price of US $375 million, subject to customary adjustments. This acquisition forms a central part of a broader eight-year strategic partnership between Wipro and Olam Group, which is expected to exceed US $1 billion in total value.

Mindsprint, established in 2007, provides digital technology services to the food and agriculture sectors. The company employs over 3,200 people across India, Singapore, the United States, the United Kingdom, and the Middle East. Its service portfolio includes enterprise applications, data analytics, cybersecurity, and supply chain technology. Mindsprint reported US $135.6 million in revenue for the 2025 calendar year, showing a steady increase from US $130.5 million in 2024.

Under the terms of the deal, Wipro will acquire 100% shareholding in Mindsprint. The company will operate as a wholly owned subsidiary of Wipro. Suresh Sundararajan, the current CEO of Mindsprint, will continue to lead the business to maintain operational continuity.

Wipro will manage Olam's end-to-end technology operations, applying AI across the entire farm-to-fork value chain to drive large-scale digital transformation.

The broader agreement includes a committed spend of US $800 million by Olam Group over the next eight years. Wipro will manage Olam's end-to-end technology operations, focusing on the "farm-to-fork" value chain. This includes farming, manufacturing, trading, and supply chain logistics. Wipro plans to apply its existing AI platforms to Olam's global operations, which serve 22,000 customers in more than 60 countries.

For Olam Group, majority-owned by Temasek Holdings, the sale aligns with its 2025 Re-organisation Plan. The company intends to monetise non-core assets and distribute net proceeds to shareholders through special dividends. Olam CEO Sunny Verghese stated that the move allows the company to focus on its core operating businesses while securing a long-term technology partner.

The acquisition requires regulatory approvals from authorities in Saudi Arabia and Australia. Both companies expect the transaction to close by June 30, 2026. Following the announcement on April 6, 2026, Wipro's shares rose 3.2% in early trading on Indian stock exchanges.

This deal represents one of Wipro's largest contract wins in recent years. And, by integrating Mindsprint's specialised IP-led solutions, such as its plantation management and procurement platforms, Wipro aims to strengthen its position in the global food and agri-business market. The partnership provides Wipro with long-term revenue visibility and adds deep domain expertise in supply chain technology. ■

Why AI velocity is outrunning human control?

AI is moving faster than humans can govern: Inside the widening velocity gap, where weak strategy, blurred accountability, and workforce fears are undermining enterprise control

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

According to a report named "Humans at the Helm of AI" commissioned by Altimerik, the global enterprise is currently navigating the AI Velocity Gap, a widening disparity between the rapid deployment of autonomous systems and the slow evolution of the human systems required to govern them. While AI is increasingly embedded in workflows, true institutionalised maturity remains rare, with only 13% of enterprises managing AI as a standardised capability.

This imbalance has created a "hollow loop" in governance; although 53% of organisations rely on human-in-the-loop processes as their primary trust mechanism, a staggering 82% of those humans lack the visibility to interrogate the reasoning behind the outputs they approve. Furthermore, when human judgment and AI recommendations conflict, only 26% of organisations have a clear rule that human judgment prevails, leaving critical decisions to be negotiated without documented authority.

The strategic vacuum and misplaced accountability

The strategic foundation for AI is thin, with only 14% of Global 2000 organisations possessing a clear strategy with defined goals. In this vacuum, 52% of leaders have defaulted to cost reduction as their top driver, a narrative that requires little vision but often results in architectures optimised for elimination rather than future revenue growth.

Accountability is also structurally misplaced; day-to-day responsibility typically rests within technical functions, yet CEOs are three times more likely to be involved only after an incident or failure occurs. This accountability without borders extends to external partners; while 83% of firms depend on partners for speed, 80% admit that responsibility is unclear when a partner's AI makes a mistake.

The culture of fear

A profound culture of fear is currently stalling genuine adoption and oversight. More than half of employees cite the fear of replacement as their primary barrier to AI engagement, and 72% fear being judged if their experiments fail. Despite these high stakes, nearly 80% of the workforce receives fewer than 10 hours of AI training per year, leading to widespread "AI imposter syndrome" where 43% of employees feel they lack the knowledge to use the technology confidently.

This has resulted in a transition by drift, where 52% of organisations expect AI to reduce roles over the next three years, mostly through attrition, yet only 7% of employees feel they are in control or shaping those outcomes.

The bottom line for the C-Suite

The next decade of business will not be defined by who built the most capable models, but by who built the most capable humans to direct them. Closing the velocity gap requires a fundamental redesign of human authority before autonomy is extended. Leaders must move beyond feeling governed and start asking the hard question: What does the human at the helm actually have the authority, visibility, and accountability to do?. ■

CIO PLAYBOOK 2026

SET CONTEXT. SHAPE DECISION. WIN BIG.

The CIO Playbook 2026, held in Varanasi, brought together India's top 100 CIOs to redefine enterprise strategy and advance AI, data, and security agendas for the next decade of transformation.

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

The CIO Playbook 2026

The CIO Playbook 2026 in Varanasi brings India's top 100 CIOs together to redefine enterprise strategy, advancing AI, data, and security agendas for the next decade of transformation

The conference opened with a unique cultural-tech fusion. **Rakesh Bedi, the iconic stage and film actor,** joined **Sachin Mhashilkar, Sr. Director, ET Edge,** for a captivating dialogue on **"Reimagining Films with AI,"** exploring how generative intelligence is reshaping the dramatic arts.

The CIO Playbook 2026, held from March 5-7, 2026 in Varanasi, brought together the nation's foremost technology leaders at the DoubleTree by Hilton in Varanasi. Under the vision of redefining the enterprise for an AI-driven decade, the summit served as a strategic crucible for over 100 CIOs to transition from digital adopters to transformation leaders.

The forum was designed to move beyond theoretical hype and surface the granular, strategic challenges of implementing autonomous intelligence at an enterprise level.

The choice of Varanasi; a city defined by its timeless heritage-provided a deliberate contrast to cutting-edge discussions on agentic workplace transformation and hybrid cloud architectures, emphasizing that true transformation requires balancing rapid innovation with foundational stability.

Throughout the three-day residency, attendees engaged in a rigorous agenda of "Interactive Technology Sessions," designed specifically for C-suite leaders to co-creat practical action frameworks through real-time brainstorming and rapid roundtable formats. From deep-dives into AI-ready infrastructure and energy optimization at scale to exploring the ethical guardrails of government and risk, the CIO Playbook functioned as a definitive roadmap for leaders tasked with architecting the next decade of corporate intelligence. ■

The first day balanced spiritual heritage with cutting-edge technical strategy. Heritage walking tour to the **Kashi Vishwanath temple**.

The formal sessions commenced with a research presentation on **"CIO Challenges and Priorities for 2026"** by **Jatinder Singh, Chief Editor, CIO&Leader, ET Edge.**

Ramesh Narayanswamy, Group President – Digital and Data Intelligence, Hinduja Group, delivered a masterclass on **"An AI Agent Playbook for Tech Leaders,"** providing a blueprint for autonomous enterprise operations.

Leaders participated in rotating workshops to deep-dive into architectural challenges.

Seema Ambastha, CEO, L&T Vyoma, and **Rohit Adnani, AI Security Architect, F5,** led interactive technology sessions on **"AI-Ready Infrastructure: Hybrid Cloud & Edge Computing Architecture"** and **"Securing the AI-Enabled Enterprise: Governance, Risk & Guardrails,"** respectively.

Apurva Jain, National Lead-Data Security Sales and Strategy, Forcepoint explored the Intelligent Data Enterprise, focusing on accelerating transformation without losing compliance control.

Robin Roy, Director, Delta addressed physical bottlenecks, focusing on Power, Cooling, and Energy Optimization at Scale.

Panel Discussion: AI at Scale: Is Your Architecture the Bottleneck (From L to R) Giridhar R, Editorial Director, ET Edge; Aashish Kshetry, CIO & VP – IT, Asian Paints; Nikhil Bhushan, CTO, Tata Starbucks; Ninad Raje, Group CIO, Times Group; Ranjitha R, Head of Engineering, PhonePe; Seema Ambastha, CEO, L&T Vyoma; and Prakash Dharmani, CIO, Okaya Power.

The day concluded on the waters of the Ganges with a luxury river cruise to witness the **Ganga Aarti.**

Soulful Shehnai concert by **Nasir Abbas Bismillah Khan**

Day 2 of the conference set the stage as **Pratik Pal, Advisor – Data and AI, Chairman's Office, Tata Group,** underscored governance, scale, and sustainable impact in his keynote, stressing that **AI must drive long-term value—not short-term hype.**

Sahin Mhashilkar, Sr. Director, ET Edge felicitated **Haji Ramzan Ali, a National Award-winning weaver,** for bridging the gap between ancient artisanal skill and modern recognition.

Panel Discussion: The CIO as Chief Transformation Officer: Redefining the Role, the Mandate, and Metrics (From L to R) Jatinder Singh, Chief Editor, ET Edge; Gaurav Duggal, CIO, Jio Platforms; Global Consulting Practice, CIO Advisory, TCS; Mayank Sharma, Chief AI & Innovation Officer, IIFL Finance; Chitti Babu, Group CIO, Aurobindo Pharma; Narrottam Sharma, CIO, Jubilant FoodWorks; Anil Shankar, Global CIO, Tata International; and Dipu K.V., Senior President, Bajaj General Insurance.

The summit wrapped with a visit to the **UNESCO World Heritage site at Sarnath.**

Is architecture the weight AI cannot carry?

Scaling AI requires a rigorous focus on data and infrastructure. To bridge the gap from pilot to production, organisations must fix data accessibility, navigate exponential scaling costs, and rethink legacy philosophies.

(L to R): **Giridhar R, Editorial Director, ET Edge; Aashish Kshetry, CIO & VP – IT, Asian Paints; Ninad Raje, Group CIO, Times Group; Nikhil Bhushan, CTO, Tata Starbucks; Ranjitha R, Head of Engineering, PhonePe; Prakash Dharmani, CIO, Okaya Power;** and **Seema Ambastha, CEO, L&T Vyoma.**

The transition of Artificial Intelligence from an experimental "innovation sandbox" to a permanent fixture on the boardroom agenda has been swift, yet the path to integration remains fraught with structural tension. At the recent CIO Playbook Conference 2026, in a compelling panel discussion titled "AI at Scale: Is Architecture the Bottleneck?", top technology leaders gathered to deconstruct a brewing crisis in the enterprise: the realisation that while models are evolving at breakneck speed, the underlying architecture may be the very bottleneck stifling its potential.

The panelists for the session were Aashish Kshetry, CIO & VP at Asian Paints; Ninad Raje, Group CIO, Times Group; Nikhil Bhushan, CTO, Tata Starbucks; Ranjitha R., Head of Engineering, PhonePe; Prakash Dharmani, CIO, Okaya Power; and Seema Ambastha, CEO, L&T Vyoma. The session was moderated by Giridhar R, Editorial Director, ET Edge.

The myth of model failure

For years, the industry narrative suggested that if an AI project failed, the algorithm was to blame. The consensus among the CIO Playbook panelists, however, suggests otherwise. "If you have a Ferrari... but you are driving the Ferrari on Mumbai pothole roads, it's not going to work out; it is going to crack open," remarked Ninad Raje, Group CIO at Times Group. This vivid analogy underscores the panel's primary thesis: AI failures are rarely model failures; they are infrastructure and data failures.

The "garbage in, garbage out" adage has never been more relevant. Panelists emphasised that readiness is not defined by the number of GPUs in a rack, but by the maturity of Master Data Management (MDM), cataloging, and meta-tagging. Without decoupling data from core legacy applications, enterprises find themselves force-feeding sophisticated AI "kids" a diet of fragmented, unreliable information.

Legacy: The three-year expiry date

One of the most provocative insights from the discussion was the recalibration of "legacy." In the era of Generative AI, panelists suggested a radical new timeline: "In terms of AI... three years is legacy".

This creates a fundamental dilemma for organisations running on decades-old systems. The panel debated the "rip and replace" versus "surround and extend" strategies. While ripping out monolithic systems is often commercially impossible—especially in sectors like manufacturing where CAPEX-heavy machinery has long lifespans—the consensus leaned toward a use-case-driven hybrid approach. For predictive maintenance or real-time recommendation engines, "surrounding" legacy systems with AI extensions is a pragmatic first step. However, for true competitive advantage, the goal must be "AI-native" architecture where intelligence is the core, not a decorative top layer.

The exponential cost of scale

If the pilot is the honeymoon, scaling is the reality check. The panel identified "The Elephant in the Room": Cost. Scaling an AI-centric application is not a linear expense; it is exponential. The money spent on scaling an application that is not AI-centric might be 'X', but for an AI-centric one, it can become '10X'.

Beyond the sheer cost of tokens and GPUs, technical hurdles like "concept drift" and "data drift" become amplified at scale. A model trained on data from a single region like Karnataka may buckle when faced with the heterogeneity of the entire Indian market. This variability requires a level of governance and continuous monitoring that many enterprise architectures are simply not yet equipped to handle.

Sovereignty and the power play

The discussion took a macroeconomic turn regarding India's "compute advantage." With estimates suggesting India will need hundreds of thousands of GPUs by 2030, the industry is witnessing a massive infusion of capital—upwards of $12 billion from single enterprises—to build sovereign compute capacity.

Yet, hardware is only half the battle. The looming shadow over India's AI ambition is power and transmission. While renewable energy generation is soaring, the transmission infrastructure and data centre efficiency (PUE) remain critical bottlenecks. The move toward liquid cooling and direct-to-chip technology is no longer an option; it is a necessity for the high-density workloads AI demands.

Where does AI live?

The final architectural debate centred on location. The verdict? AI should live wherever the data resides. Whether it is edge computing for QSR (Quick Service Restaurants) to minimise latency, or on-premises setups for sensitive shop-floor telemetry, moving petabytes of data for inference is a "project killer".

As the panel concluded, the message for the modern CIO was clear: AI success is a game of 90% preparation and 10% execution. To win, one must run several parallel experiments just to find the one that sticks. But more importantly, one must ensure that the "roads" of the enterprise architecture are ready for the "Ferrari" of AI. ■

Data, not models, is the true frontier of AI at scale

Ashish Kshetry, CIO & VP, Asian Paints

The prevailing boardroom narrative often suggests that AI failures are primarily algorithmic—that a "better" model is the silver bullet. However, the reality on the ground is that data readiness, not the model, is the single biggest hurdle to enterprise AI. To move beyond the sandbox, organisations must decouple data from core applications, focusing on master data management, cataloging, and meta-tagging to make information truly useful for subsequent modeling.

In a heavy-industry environment like manufacturing, we cannot simply "rip and replace" high-CAPEX legacy machinery. Scaling AI requires a hybrid architectural approach. For instance, achieving predictive maintenance on a shop floor often requires an eight-month journey of working with OEMs to replace PLCs and establish static IPs for real-time data streaming. You don't need to re-architect everything at once, but for critical use cases, AI must eventually move from a "top layer" addition to being natively embedded at the core.

While AI projects for internal productivity are manageable, the "elephant in the room" is cost when scaling to B2C applications. Exposing AI functionalities to the vast population of India brings exponential GPU and token costs that can cripple a project if not planned carefully. Where should these workloads live? It is not one-size-fits-all. While public clouds are ideal for non-sensitive data modeling, mission-critical telemetry must stay at the edge. On the shop floor, the latency involved in sending data to the cloud and back will simply "kill" the use case. Ultimately, the decision must be driven by a trifecta of data sensitivity, sovereignty, and the physical requirements of the machine. ■

Why infrastructure is the true AI bottleneck

Nikhil Bhushan, CTO, Tata Starbucks

For organisations to achieve success with AI, a perfect harmony must exist between data, infrastructure, and compute. While models are constantly evolving and far less costly to update, the real "expensive fix" lies in infrastructure that cannot keep pace or data sets that lead to flawed results. Scaling in a territory like India is particularly challenging because the country is essentially "many countries in one," making it impossible to rely on a one-size-fits-all build across diverse micro-markets. Furthermore, leaders must realise that as they scale, infrastructure requirements grow exponentially rather than linearly.

This physical reality is complicated by a power-deficit environment heavily dependent on thermal generation, raising critical concerns about whether generation can keep pace with massive new data centres. At this scale, every enterprise must decide if an initiative brings a competitive advantage that is easy to replicate or a proprietary differentiation that can be owned long enough to recoup the investment before the competition inevitably catches up. For businesses in the Quick Service Restaurant (QSR) sector, the solution for these workloads is often at the edge. To service a customer at the point of sale with minimal latency, compute must be decentralised to ensure the speed and responsiveness modern consumers demand. ■

Beyond the myth of the architectural bottleneck

Ranjita R, Head of Engineering, PhonePe

The industry often laments that architecture is the primary bottleneck for AI, but I believe this is a myth. AI is essentially like a child: garbage in, garbage out. The real issue is that organisations try to force-feed AI into existing systems riddled with anti-patterns and legacy philosophies. To succeed, you must first correct your architecture according to software first principles; then, AI will fit in naturally.

Legacy is not just old software; it is a different philosophy of thought. If you are running batch processing systems that ingest data every six to eight hours, you cannot expect near real-time AI recommendations. It is about the use case. While a batch system might still work for detecting daily anomalies, a streaming use case requires a complete replacement of that philosophical approach to work effectively.

At scale, the "drift" problem in AI—both concept and data drift—becomes amplified. A model might perform beautifully in a closed, controlled environment with limited data. However, as you scale from a single region like Karnataka to the entirety of India, data heterogeneity increases, and the initial guardrails often fail. The variables simply change when you move to a national scale. Determining where workloads should live requires asking what you are willing to compromise on: latency, cost, or security. There is no one-size-fits-all answer. My philosophy is to start in the cloud for the initial experimentation phases. Once a model is stable, working, and showing predictable loads, you should insource it. This hybrid approach allows you to manage costs and ensure optimal utilisation while maintaining the flexibility to handle petabytes of data securely. ■

The geopolitical race for sovereign compute and the cost of AI business

Seema Ambastha, CEO, L&T Vyoma

The most critical hurdle to AI readiness is not the model, but data accessibility. While enterprises believe they are data-rich, making that data accessible to AI is the "one fix" upon which everything else rolls. Furthermore, scaling AI is not a standard linear progression; while scaling conventional CPU applications costs 'X', AI-centric scaling costs '10X'. This cost, coupled with data accessibility, is why many pilots that look fantastic in controlled environments fail to deliver at scale.

We are currently in a geopolitical race driven by compute. To ensure India is not left behind, the industry must move beyond small-scale GPU grants. L&T is pledging billions—starting with $2 billion and reaching up to $12 billion over five years—to create a national compute advantage of 50,000 to 100,000 GPUs. We estimate India will require 400,000 GPUs by 2030, yet the current capacity is only around 25,000. To manage the "brunt" of these costs, we are looking at five-year amortisation models to give enterprises a better advantage. India possesses ample power generation, exceeding solar and wind projections. The bottleneck is transmission and substation upgrades. We must also focus on efficiency, driving Data Centre Power Usage Effectiveness (PUE) from 2.0 down to 1.1 using liquid cooling and direct-to-chip technologies.

Finally, we must correct a fundamental misunderstanding: 95% of AI workloads do not require moving data during inference. Training is a one-time activity; post-training, you need sturdy inference engines, not mass data movement. If your model relies on moving petabytes of data for inference, it isn't AI—it's just a transaction. ■

Why hardware isn't AI readiness

Ninad Raje, Group CIO, Times Group

A dangerous myth persists in modern boardrooms: that purchasing GPUs or a "fancy" platform equates to AI readiness. It is a fundamental misconception. Think of it like buying a Ferrari to navigate the pothole-ridden streets of Mumbai; the hardware is elite, but the infrastructure will cause it to crack. True readiness is about interoperability and clean, reliable, current, and accessible data that allows AI tools to generate fact-based insights for knowledgeable growth.

Scaling AI requires a rigorous approach to data discovery. Within an organisation dealing with petabytes of data monthly, you must first identify what is relevant. This involves a three-tier process: identification, relevancy, and governance. Without strict governance, unmanaged data creates more problems than it solves. You must extrapolate this refined data to match specific organisational needs and business plans. I apply a simple, brutal filter to any AI-centric project: Will this reduce costs or increase revenue by tomorrow? If the answer is no, we don't proceed; we go back to the experimentation board. Success at scale also requires a high volume of trial and error. Based on our experience, you must run nine to ten parallel experimentation projects to yield just one or two successful outcomes. It is those few successes that ultimately satisfy the board and the customers.

For AI to truly scale, we must move away from the obsession with hybrid or edge mechanisms as defaults. AI platforms must be as close as possible to the data source. Moving petabytes of data from storage to a distant compute location is simply not feasible. It downgrades performance, increases costs exponentially, and destroys the user experience. If the data is on-prem, the AI must be there; if it's in the cloud, the AI must follow. Anything else is a recipe for failure. ■

Bridging the gap—Data readiness and infrastructure sovereignty

Prakash Dharmani, CIO, Okaya Power

The industry often fixates on model failures, but the true culprit is data readiness. AI functionality depends on how well data is decoupled from core applications and integrated into a broader ecosystem through master data management and robust meta-tagging. It's not just a myth; it's a practical challenge of infrastructure readiness. In a manufacturing context, we cannot simply "rip and replace" expensive legacy machinery. Instead, we must adopt hybrid architectures, such as replacing PLCs with static IPs to enable real-time data streaming for predictive maintenance—a process that can take months of coordination with OEMs and SIs. .

Scaling these solutions brings us to the "elephant in the room": cost. While internal projects like supply chain optimisation remain manageable, scaling to B2C applications for the general population introduces exponential GPU and token costs. This physical reality is further complicated by national energy challenges. While distributed solar and decentralised power generation are promising, the immediate reality of a power-deficit environment dependent on thermal energy could pose a significant challenge for new data centres by 2030.

Strategically, the decision of where to house these workloads must be driven by data sensitivity, sovereignty, and latency. For non-sensitive modeling, the public cloud is efficient. However, for sensitive data or mission-critical shop floor telemetry, the workload must remain at the edge. The latency involved in round-tripping data to the cloud will simply kill the application. Ultimately, you must decide based on the specific end-use and the physical requirements of the environment. ■

Sumedha Chakraborty
India Head,
Google Workspace

Redefining business outcomes through workplace transformation

The next phase of AI is no longer assistance—it is enterprise action that reshapes how work is executed and business outcomes are delivered.

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

For the modern enterprise, 2026 marks a fundamental pivot in the AI journey. We have moved past the era of the "chatbot"—where success was measured by prompt proficiency—into an era focused on Agentic Transformation. For CIOs, the mandate has evolved from merely deploying tools to orchestrating a cohesive ecosystem that bridges the gap between digital collaboration and core business execution.

The shift from assistance to agency

The first wave of generative AI provided individual assistance—summarizing documents or drafting emails. However, true business-wide transformation occurs when AI moves from a passive sidebar to an active participant in business processes. This Agentic approach allows organizations to move beyond task-level efficiency toward the automation of complex, multi-step workflows.

By fostering a unified environment where AI-powered assistance is embedded into the fabric of daily work, enterprises can break down traditional data silos. This integration ensures that the "Future of Work" is not just about faster communication, but about higher-velocity decision-making grounded in the organization's unique data and operational workload.

Driving value through cultural and operational scaling

The roadmap to 2026 is defined by moving from experimentation to production at scale. Successful organizations are focusing on three core strategic pillars:

- **Redefining productivity:** Shifting the workforce from routine transactional execution to high-value strategic direction. This requires a rethink of how we measure output in an era where AI handles the operational heavy lifting.
- **Bridging the strategy-execution gap:** Moving beyond isolated AI pilots toward connected systems that can reason across departments—from Finance to Supply Chain—ensuring that insights lead directly to autonomous or semi-autonomous action.
- **Governance by Design:** Building a foundation of trust where AI is safe, governed, and transparent. The competitive advantage in 2026 will belong to those who treat data sovereignty and responsible AI not as constraints, but as the primary enablers of scale.

The human-centric advantage

Technology is only half the equation. The defining characteristic of a leading enterprise in 2026 is its focus on human-AI collaboration. As we move into this next phase, the goal is to empower people to focus on the complex challenges that require human judgment, empathy, and creativity.

Sources & Strategic References

- Google Cloud Blog: 5 ways AI agents will transform the way we work in 2026
- Transform with Google Cloud: Beyond the pilot: Five hard-won lessons from Google Cloud's AI transformation strategy
- Google Cloud Product Strategy: Gemini Enterprise: The new front door for AI in the workplace
- Google Cloud Leadership: The path to an AI-powered organization ■

Seema Ambastha
Chief Executive,
Larsen & Toubro-Vyoma

The AI problem isn't models, it's your process

AI at scale will be won on infrastructure, not models—how sovereign, hybrid, and edge-first architectures will define the next enterprise frontier.

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

For the past decade, we treated AI as an experimental capability- tucked away in innovation labs or deployed in isolated, low-risk use cases. Today, that paradigm is extinct. AI has become a board-level imperative, tasked with driving structural transformation across the global economy. Yet, as we push these ambitions forward, many of my peers in the C-suite are confronting a harsh reality: The primary barrier to AI at scale is no longer the quality of the model, but the integrity of the infrastructure.

Most enterprise IT estates were built for transactional stability- ERP systems, reliable databases, and predictable, structured workflows. AI, however, demands a radical departure. It requires massive GPU-accelerated clusters, high-throughput networking, and real-time inference capabilities that our legacy foundations were never designed to sustain.

The reality is clear: AI will not be constrained by model innovation. It will be constrained by infrastructure readiness. Enterprises that architect their AI infrastructure today will define the digital economy of tomorrow.

Is your IT "AI-ready"? The cost of architectural friction

We are witnessing a common bottleneck across the enterprise landscape: many organizations remain trapped in the "pilot phase." This is not a failure of strategy; it is a reflection of architectural readiness. For the global CIO and CTO, the friction is threefold:

- **The legacy debt:** Attempting to run GPU-intensive, high-velocity AI pipelines on infrastructure optimized for 2010s-era transactional stability. This creates a performance ceiling that no algorithmic fine-tuning can bypass.
- **The data paradox:** Organizations have vast volumes of data, but it is often fragmented, ungoverned, and trapped in silos. Without a unified data-fabric that spans from the core to the edge, the AI agents are fed with "dark data" that lacks the lineage, provenance and above all – factual relevance to the enterprise that is required for high-stakes decision-making.
- **The talent-infrastructure gap:** The complexity of managing heterogeneous, GPU-enabled environments is outstripping the current capacity of internal teams. With the new class of "AI-native" management needs, the emphasis on infrastructure as code couldn't be stronger.

Sovereign AI: Compliance as an architectural moat

Perhaps the most significant shift in the global mandate is the rise of Sovereign AI. As data residency, compute jurisdiction, and supply chain risks enter the C-suite vocabulary, leaders are realizing

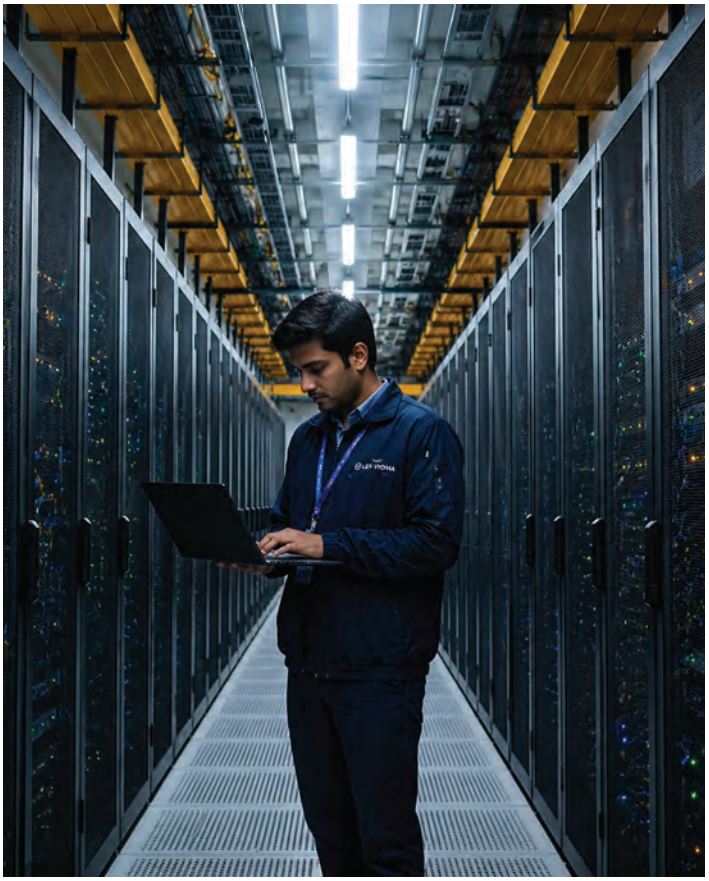

"Enterprises that architect their AI infrastructure today will define the digital economy of tomorrow."

that relying solely on public cloud hyperscalers could become a strategic liability.

For the regulated enterprise, Sovereign AI is not just about compliance- it is about sovereignty over your most proprietary intelligence. Modern regulatory frameworks globally are introducing risk-tiered obligations, which are not just bureaucratic hurdles; they are defining the new requirements for data lineage, consent management, and auditability.

At Larsen & Toubro-Vyoma, we argue that compliance can no longer be "layered on" after deployment. It must be architected into the infrastructure design. We are helping organizations build private and community clouds that ensure data localization, model independence, and verifiable AI incident response. By baking governance into the silicon and the software-defined fabric, we allow companies to innovate rapidly without triggering the "regulatory brakes" that often stall AI deployments.

The new paradigm: Hybrid cloud + The distributed edge

To achieve enterprise scale, we must move toward a distributed architecture. AI workloads are inherently hybrid: training belongs in the hyperscale environment, while inference- particularly for manufacturing, autonomous systems, or smart city infrastructure- must reside at the "Edge," where latency is the enemy of performance.

Larsen & Toubro-Vyoma is uniquely positioned at the intersection of digital engineering and heavy physical infrastructure. We aren't just building data centers; we are architecting a Hyperconnected Intelligence Ecosystem. Our approach focuses on four critical pillars:

- **GPU-As-A-Service (GaaS):** Providing the raw, sustainable compute power specifically tuned for high-density GenAI and RAG pipelines.
- **Edge-to-Core Fabric:** A unified MLOps framework that allows a model to be trained in a massive central facility and deployed seamlessly to an edge device on a factory floor or telecom network, thousands of miles away.
- **Sustainability by Design:** Leveraging Larsen & Toubro's heritage in industrial engineering to build power-dense, green-energy-backed clusters that can handle the massive heat dissipation requirements of next-gen AI hardware.
- **Advisory-Led Deployment:** Recognizing that infrastructure transformation is as much about people and process as it is about hardware, we provide the consulting and management expertise to guide the transition from monoliths to API-enabled AI pipelines.

The path forward: Beyond the pilot

The successful enterprise of the next decade will not necessarily be the one with the most advanced LLMs. It will be the one with the most resilient, sovereign, and distributed infrastructure.

We are moving away from the era of monolithic IT into the era of continuous, agentic AI deployment. As we look ahead, the core question for every C-suite leader is: Which of your current systems are actively blocking your scale, and what is the cost of deferring their modernization? At Larsen & Toubro-Vyoma, we are bridging this divide. We are not just building the physical data center; we are engineering the intelligent foundation upon which the future of the enterprise will be built.

The future isn't just intelligent; it's engineered. And it begins with the foundation you lay today. ■

Vaishali B. Shivshankar
Head of Marketing,
India & SAARC, F5

Navigating the AI frontier: strategic imperatives for secure and compliant innovation

Enterprise AI goes beyond speed, demanding security, trust, and control

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

The year 2026 finds enterprise leaders—CIOs, CISOs, CDOs, CROs, and CDIOs—at a critical juncture. Generative AI, once a nascent technology, has matured into a transformative force, reshaping business models, enhancing customer experiences, and driving unprecedented operational efficiencies. Yet, this rapid evolution also ushers in a new landscape of complex risks that demand a proactive, integrated security strategy. As organizations race to harness AI's potential, the imperative is clear: innovation must be built on a foundation of robust security and unwavering compliance.

The unprecedented risk profile of generative AI

Traditional security paradigms, honed over decades to protect applications and data, are proving insufficient against the unique vulnerabilities of AI models. The attack surface has expanded dramatically, introducing novel threats such as prompt injections, jailbreaks, system prompt exposures, and sophisticated obfuscation techniques designed to bypass security controls. Data privacy concerns intensify with the potential for large language models (LLMs) to inadvertently expose personally identifiable information (PII) or generate toxic and biased content.

For the CISO, this means a constant battle against evolving adversarial tactics targeting the logic and data integrity of AI systems. The CRO, meanwhile, must grapple with broader enterprise risk implications—from reputational damage and financial losses to regulatory penalties driven by legislation such as the EU AI Act, which by March 2026 is enforcing significant compliance mandates across global operations. CDOs face the challenge of ensuring data governance and model integrity, while CIOs and CDIOs must balance rapid AI adoption with the need for secure, trustworthy deployments that deliver tangible business value without compromising the enterprise.

F5's holistic vision: securing the entire GenAI project lifecycle

Recognizing this multifaceted challenge, F5 offers a comprehensive AI security platform designed to empower leaders with control without compromise. It integrates seamlessly across the entire GenAI project lifecycle—from model selection and solution development to production inference—ensuring security is embedded at every stage. This platform approach addresses the concerns of every C-suite stakeholder, providing tools for proactive risk identification, continuous threat validation, and active defense.

"AI innovation must be built on a foundation of robust security and unwavering compliance."

At the core of F5's offering are three synergistic pillars:

F5 AI Red Team: proactive threat intelligence and model hardening:
For the CISO and CRO, understanding the inherent risk profile of an AI model is paramount. The F5 AI Red Team acts as an advanced threat simulation engine, conducting continuous adversarial testing to stress-test models against evolving threats. It goes beyond simple vulnerability scanning, analyzing the agentic chain of thought to plan and execute multi-turn attacks based on custom intents. This capability identifies model risks—quantified by a CASI score—and provides detailed analysis and remediation recommendations. Organizations can review model scores, export findings via PDF or API, and analyze specific test data. The F5 Labs CASI Leaderboard offers an industry benchmark, aiding CDOs and CIOs in selecting robust and secure models. This proactive posture enables enterprises to "start secure" by aligning controls with their risk appetite, reducing downstream exposure.

F5 AI guardrails: enforcing security and compliance in real time:
As AI models move from development to production, consistent enforcement of security policies becomes critical. F5 AI Guardrails delivers active defense by blocking attacks at inference and enforcing controls during development. For CDOs and CIOs, this ensures data integrity and responsible AI usage. For CISOs and CROs, it strengthens compliance and reduces attack vectors. Key capabilities include:

- AI security guardrails: protection against prompt injections, jailbreaks, system prompt exposures, and obfuscation/bypass attempts
- Topic moderation guardrails: prevention of inappropriate outputs (e.g., financial or health advice)
- Compliance guardrails: alignment with regulations such as the EU AI Act
- PII masking/blocking: automatic detection and redaction of sensitive data in prompts and responses

Organizations can calibrate these controls to match their policies and innovation strategies or build custom guardrails using natural language prompts or datasets.

AI observability: enterprise-wide visibility and control
Effective AI security and performance require end-to-end visibility. F5's AI observability provides real-time monitoring of AI inference across the enterprise. This enables CIOs and CDIOs to track model performance, security events, and compliance continuously. Integrated with DevOps automation through API and Python SDKs, it embeds security insights directly into development and deployment pipelines, enabling a secure and iterative lifecycle.

Seamless integration for robust protection

F5 AI Security supports flexible integration architectures—both out-of-band and inline—to suit diverse environments. In an out-of-band setup, applications send prompts to the F5 AI Guardrails API for assessment, receiving a pass/flag response before execution. This provides a lightweight, non-intrusive layer of security.

In an inline architecture, F5 AI Guardrails functions as an OpenAI-compatible chat completions API proxy. Prompts are assessed, routed to model inference, and responses are evaluated—and, if necessary, blocked or redacted—before reaching the user. This approach ensures real-time, embedded security enforcement within AI interactions.

The F5 application delivery and security platform: a foundation for the future

These AI security capabilities are built on the broader F5 Application Delivery and Security Platform, offering a unified approach to protecting traditional, modern, and AI-powered applications. The platform enables control without compromise, delivering application security and performance across environments and form factors. It also fosters XOps collaboration across NetOps, SecOps, and DevOps, reducing blind spots and strengthening enterprise resilience.

As of March 2026, the strategic imperative for CIOs, CISOs, CDOs, CROs, and CDIOs is clear: embrace AI with confidence, ensuring that security and compliance are foundational—not optional. F5's integrated AI security platform provides the tools and framework needed to navigate this rapidly evolving frontier, enabling organizations to realize AI's full potential securely and responsibly. ■

AI doesn't need better models. Enterprises need better processes

Ramesh Narayanswamy, Group President – Digital and Data Intelligence, Hinduja Group, on why AI success hinges more on mindset, process redesign, and integration than on advanced models or scarce talent.

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

At most enterprise gatherings on artificial intelligence, the room usually fills with familiar anxieties—lack of talent, too much uncertainty, unclear ROI, and a technology curve moving faster than institutions can absorb.

At the recent CIO&Leader conference 'The CIO Playbook' held in Varanasi from 5th to 7th March, Ramesh Narayanswamy, Group President – Digital and Data Intelligence, Hinduja group delivered a compelling keynote address on "An AI Agent Playbook for Tech Leaders", setting the tone of the conference.

In his insightful keynote session, Ramesh chose not to soothe any of these concerns. Instead, he challenged one of the central myths shaping the AI era: that expertise lies elsewhere. The real work

of AI transformation will not be outsourced to consultants, absorbed by vendors, or solved by buying another layer of software. It will be led—or lost—by the people already inside the organisation.

Beneath the candor was a sharper warning for CIOs, CTOs, and boards alike: in the AI economy, enterprises that merely automate inefficiencies will deepen them; those that rethink process, capability, and decision architecture will define the next operating model.

Collapse of old hierarchy

For all the noise around AI maturity, most enterprise practitioners have only a narrow window of real experience with generative AI. That reality, he suggested, should be liberating rather than intimidating. It means the old hierarchy of expertise is weaker than many assume.

Consultants, system integrators, and OEMs are often learning at the same pace as enterprises themselves. The winners in AI will not necessarily be those with the most polished vocabulary, but those willing to learn in public, build in context, and act without waiting for perfect certainty.

Swamy insisted that the real obstacle to AI adoption is not technology—it is process. For years, enterprise IT has operated within a familiar division of labour: business defines the process, and technology builds systems around it. AI breaks that arrangement.

In his view, many existing processes are simply not fit for an AI-enabled enterprise. Unless leaders rethink workflows end-to-end, these projects stall, underperform, or collapse into experimentation theatre.

Process, not models, is the real battlefield

He described this as the first barrier technology leaders must remove from their thinking. Too often, organisations begin with the data they already have and attempt to build around it. AI flips that sequence.

The starting point is no longer, "What data do we have?" It is, "What process are we trying to reinvent?" Only then does the data question become meaningful.

AI is often discussed as a capability layer. In practice, it is increasingly a trigger for institutional redesign. The organisations creating real advantage are not just automating steps; they are questioning whether the process itself deserves to survive.

That demands something many technology functions have historically avoided: stepping into domains once owned by operations or finance and actively challenging how work gets done.

"If you don't redo the process end-to-end, these projects go nowhere. It doesn't matter what technology you use."

Talent Is not the constraint—Mindset is

Swamy's remarks on talent were equally pointed. Across industries, talent scarcity has become the default explanation for slow AI progress. He rejected that view almost entirely.

The talent crisis, in his telling, is often a convenient excuse, partly because leaders themselves are still climbing the learning curve and partly because "AI talent" has been over-romanticised. Outside frontier labs, most enterprises do not need armies of model scientists. They need teams willing to learn quickly, experiment responsibly, and work across process, data, and systems.

His point was not that skill does not matter, but that enterprises may be looking for the wrong kind of scarcity.

AI is not a model problem

Swamy's account of building an AI platform from scratch cut through much of the abstraction that

"We always start with data—what we have and what we can build. But AI flips it. First define the process, then find or create the data."

dominates AI conversations. His team built not for show, but to understand the mechanics, and to evaluate what vendors were selling.

The outcome was instructive. They concluded they did not need to develop their own LLMs, obsess over fine-tuning, or anchor their strategy on the most glamorous parts of the stack.

What emerged instead was a harder truth: enterprise AI is largely an integration challenge.

Where does data live? How clean is it? How is it labelled? How reliably can it move across systems? How well do process logic, architecture, and decision-making align?

AI has not removed these problems—it has exposed them more sharply.

The ROI Illusion

Then came the issue that haunts nearly every boardroom conversation on AI: return on investment.

Swamy addressed it with realism. Technology has rarely been funded through perfect upfront ROI logic, and AI should not be subjected to a false precision that other major transformations were never required to meet.

The deeper problem, he argued, is that enterprises still measure AI using budgeting logics built for a slower era. If a project planned for six months is completed in three, organisations often lack mechanisms to redeploy capital, teams, or momentum dynamically.

Traditional annualised planning distorts the value of speed. That means the ROI challenge may be less about AI itself and more about the rigidity of surrounding processes, especially finance.

Uncomfortable truth about enterprise AI

The larger message of the moment is unmistakable: AI is no longer a side experiment, a sandbox for innovation teams, or a future bet waiting for maturity. It is becoming the lens through which enterprise relevance, speed, resilience, and leadership will be judged.

The companies that win will not necessarily be those with the biggest budgets or earliest pilots, but those that can translate AI into operating reality; through cleaner data, sharper governance, redesigned workflows, leaner teams, and a culture willing to rethink work from first principles.

For CIOs, this is not merely a technology transition—it is a test of strategic courage.

The pilot era is ending. What comes next will separate enterprises that merely use AI from those that are rebuilt by it. ■

AI moves from hype to core strategy in India

Indian enterprises must operationalize AI rapidly or risk falling behind, says **Pratik Pal, Advisor – Data and AI, Chairman's Office, Tata Group,** in his keynote at the CIO Playbook Conference

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

Artificial intelligence has crossed the threshold from curiosity to compulsion. In boardrooms across India, the conversation has shifted decisively from What can AI do? To Why aren't we deploying it yet?

Models are improving, costs are collapsing, and capabilities once considered experimental are becoming operational. But yet, inside most enterprises, reality still lags ambition. The past year was not defined by transformation, but by pilots. And, that gap between external acceleration and internal inertia is now becoming untenable.

These themes were at the core of a keynote titled "The AI-First Enterprise: Governance, Scale, and Sustainable Impact," delivered by Pratik Pal, Advisor – Data and AI at the Chairman's Office, Tata Group at the CIO Playbook 2026. Drawing on his experience of building and scaling enterprise AI platforms across complex, large-scale environments, Pal grounded the conversation in execution rather than abstraction, focusing on what it truly takes to move AI from experimentation to enterprise-wide impact.

The tension resides at the epicentre of this moment. On one side is relentless innovation from hyperscalers. On the other hand, the slower machinery of the enterprise—legacy systems, fragmented data, compliance constraints, and processes never designed probabilistic systems.

The shift is now unmistakable; the question is no longer whether AI matters. It is whether organisations can operationalize it before the opportunity hardens into disadvantages.

Hype to use cases

Enterprise AI in 2025 was largely defined by utilities like code generation, chatbots, enterprise search, summarisation and workflow automation. These are mere productivity improvements, missing the edge of value creation.

What will redefine this year is the beginning of change in the way AI is applied. Applications that are not peripheral but; they sit close to revenue, risk, and operations.

The production problem

Even as AI moves into higher-value workflows, the transition from pilot to production remains the defining challenge. The barrier is not just technical, but they are institutional.

> **"What used to be testing, has now become eval, because AI is a black box."**
>
> **Pratik Pal,**
> Advisor – Data and AI at the Chairman's Office, Tata Group

Regulation is the first constraint. In sectors like financial services, compliance is not layered after deployment; it shapes it. A significant portion of effort goes into ensuring systems are defensible before they go live.

The second is accuracy. Enterprises are built on deterministic systems, clear inputs and predictable outputs. AI is probabilistic. That gap matters when decisions after loans, claims or compliance outcomes.

An 80–85% accuracy rate may be acceptable in consumer AI. In enterprise workflows, it is often unusable. Systems must be reliable, explainable, and auditable.

This is why evaluation—eval—is becoming central. Unlike traditional testing, eval is continuous. It measures and refines behaviour emerging from a black box. Enterprises that treat evaluation as an afterthought will struggle to scale. Those that operationalize it early may gain both control and speed.

Explainability as infrastructure

If accuracy determines trust, explainability determines governance. Enterprises must be able to justify decisions—why a loan was rejected, why a claim was denied, why a risk was flagged.

Explainability is not a feature. It is infrastructure. Without it, AI remains confined to low stakes use cases. With it, organizations can push into core workflows where decisions carry financial and reputational weight.

Why enterprises are still harder than models

The idea of the AI-first enterprise is compelling but often oversimplified. Enterprises are not blank slates. They are built on decades of systems-core platforms, ERP systems, CRM layers, spreadsheets, and deeply embedded workflows.

AI cannot simply be layered onto this complexity. It must work through it. That requires a differentiating approach, one that starts with understanding the enterprise before attempting to transform it.

Building an enterprise AI stack

The keynote's answer was clear; every organisation must build its own AI stack. Not a generic one, but one that is shaped by its systems, data, and risk environment.

The foundation is fourfold; infrastructure, AI-ready data, models and an orchestration layer.

Of these all, data is the most critical one. This is where many enterprises falter. Not because they lack models, but because their data is incomplete, unstructured, or inaccessible. Without a strong data foundation, no AI strategy will deliver outcomes.

AI for business transformation

The most important idea shared by Pal was; AI is not primarily a technology shift.

"AI to me is a business change, not a technical change."

That reframes the conversation. The real levers of value lie in process redesign, organisation structure, governance, and culture. Models will improve and commoditise. What will differentiate enterprises is how they adapt around them.

The models will keep improving. Costs will keep falling, Capability will continue to expand. What is not guaranteed is whether enterprises can translate that into transformation.

For CIOs the mandate has been made clear. AI is no longer a tool to experiment with. It is a question of how the enterprise itself must change. Those who treat it as incremental will gain efficiency. Those who treat it as structural will define the next operating model. ■

Building AI models that work for India

India's AI journey is at an inflection point. Over the last two years, we have moved from experimentation with large language models to serious conversations around domain-specific and small language models.

By **Sanjeev Azad** | editor.tech@timesgroup.com

India's AI journey is at an inflection point. Over the last two years, we have moved from experimentation with large language models to serious conversations around domain-specific and small language models (SLMs). The narrative is shifting from "Can we build with AI?" to "Can we build AI that truly works for India?

In theory, domain-specific models should outperform large, general-purpose systems. They are leaner, more cost-efficient, and better tuned to contextual nuances. In high-volume, cost-sensitive environments such as automated KYC, telecom support, claims processing, or supply chain reconciliation, SLMs can deliver faster, more predictable outcomes. But in Indian enterprise environments, theory often collides with reality.

The reusability crisis

Most Indian enterprises operate on heavily customized versions of global standards. A bank's

core processes may deviate significantly from international frameworks. A telecom provider's workflow might be shaped by years of legacy integrations and localized adaptations. This results in Data silos that are deeply proprietary and inconsistently structured.

When a domain-specific model is trained on such data, it becomes tightly coupled to that one organization's logic. Porting it to another enterprise, even within the same sector, can mean starting from scratch. This is what we call the "reusability crisis." The model isn't the problem; the foundation is.

If data structures are fragmented and workflows are bespoke, the ROI of any specialized model erodes quickly. Enterprises often underestimate the hidden cost of cleaning legacy data, restructuring taxonomies, and repeatedly fine-tuning the model. In many cases, these costs can be multiples of the model build itself.

The real differentiator, therefore, is not the intelligence of the model but the maturity of the data.

From cost per token to value of opportunity

Today, forward-looking leaders are evaluating AI in terms of the total value of opportunity. If an SLM can automate a high-frequency compliance task, reduce turnaround times, or unlock a new revenue stream, its impact must be measured against strategic outcomes, not just infrastructure savings.

However, there is an uncomfortable truth – Until data is standardized and governance is embedded at the architecture level, ROI remains a ghost metric. Pilot projects may show promise, but they struggle to scale. The companies that will succeed are not necessarily those with the most advanced models, but those disciplined enough to fix their "boring" processes before layering AI on top.

Sanjeev Azad
Global Chief Innovator & CTO (APAC) at GlobalLogic

Governance is a technical inclusion, not a policy PDF

India's regulatory environment is rapidly evolving, particularly with the implementation of the Digital Personal Data Protection Act, 2023. While policy conversations are necessary, AI governance in practice is less about high-level principles and more about data hygiene.

The real risk often lies in the interface layer. When AI tools begin to "remember" user history, personalize outputs, or integrate across internal systems, they create new exposure surfaces. In organizations where access controls are loosely defined and shared drives are poorly structured, AI becomes a powerful searchlight, surfacing information that was never intended to be discoverable at scale.

Effective governance cannot be an afterthought.

We cannot expect AI systems to respect privacy if our own internal data environments are chaotic.

Domain-specific Indian languages: The inclusion imperative

Another critical factor for AI adoption at scale in India is the intersection of domain expertise and Indian language capability, while including code-mixed usage.

Generic intelligence without domain grounding produces generic outputs. English-only systems limit adoption to a fraction of the population. For AI to be meaningful in India, it must speak the language, literally and contextually, for its users.

Initiatives such as Bhashini and emerging players like Sarvam AI demonstrate how localized models can power translation, speech recognition, and sector-specific applications across agriculture, healthcare, education, and financial inclusion.

Small language models optimized for vernacular contexts, device-light environments, and offline deployment can democratize access. They allow enterprises and government bodies to retain sensitive data locally, comply with regulatory expectations, and serve citizens in low-connectivity regions. When paired with voice and multimodal interfaces, these systems extend usability far beyond urban, English-speaking audiences.

2026 is going to be the year of pragmatism

By 2026, we will see a clear divide between organizations that "did AI" and those that made it work. The winners will not be those chasing the most fashionable models. They will be the ones who applied radical pragmatism, focusing on high-frequency, high-impact processes, embedding governance by design, and institutionalizing data discipline.

Equally important is the workforce. Employees must move beyond learning how to prompt a model. They need to understand how to validate outputs, detect bias, and supervise AI systems operating on localized data. Sustained capability and not one-off pilots will define success.

In the end, India's AI advantage will not come from building the biggest models. It will come from building the right foundations. This includes standardized data, responsible governance, domain depth, and inclusive language intelligence. Only then will domain-specific AI move from experiment to enterprise backbone. ■

AI projects fail when the business context isn't clearly defined

Vijayant Rai, Managing Director at Snowflake outlines how enterprises in India are moving from AI pilots to production, highlighting data readiness, unified platforms, and agentic AI as key enablers, while addressing ROI concerns, governance, and evolving workforce dynamics.

By **Jatinder Singh & Punam Singh** | editor.tech@timesgroup.com

Enterprise AI adoption is entering a decisive phase, where the conversation is shifting from experimentation to measurable business outcomes. While proof-of-concept projects have dominated the past few years, organisations are now under pressure to scale AI into production, demonstrate ROI, and integrate emerging paradigms such as agentic AI into core operations. This transition is particularly critical in markets like India, where digital scale, regulatory evolution, and infrastructure readiness are shaping how enterprises operationalise AI.

In this conversation, **Vijayant Rai, Managing Director for India at Snowflake,** outlines how enterprises are progressing from pilots to production, the challenges of fragmented data and unclear business context, and the growing importance of unified data platforms. He discusses the role of partnerships, the balance between sovereign AI and global innovation, and how India's scale, talent pool, and digital public infrastructure can accelerate AI adoption.

CIO&Leader: AI pilots remain stuck for many enterprises, even as the conversation shifts to agentic AI. How do you see this evolution, and how is Snowflake addressing the move from pilot to production?

VIJAYANT RAI: If you look at it, we are at a point where we have just released our end-goal results about 10 days ago. We finish our year in January, so for FY26, we released our financial results. We reported roughly US $4.4 billion in product revenue, which reflects actual consumption on the platform. Another important point is that we now have 13,300 customers. What stood out was that 9,100 of those customers have adopted AI on the Snowflake platform.

We are a fully managed AI Data Cloud. As customers start using AI features on top of their data estates on Snowflake, we are able to see that usage. When 9,100 out of 13,300 customers are using AI, it clearly indicates a strong movement in adoption.

There are different phases, as you mentioned—experimentation, proof of concepts, and then production. What we are seeing over the last year or year and a half is that we have moved from POCs to production. We are now seeing real-life production use cases globally and in India. Many organisations are able to see production impact and benefits. Internally at Snowflake as well, this is happening at scale.

CIO&Leader: Is adoption homogeneous across organisations? Is everyone able to scale at the same level?

VIJAYANT RAI: Not really. Some challenges exist, and they vary across organisations. There are cases where organisations try to fit generative AI into existing processes and systems. That often becomes challenging and creates more complexity. There are also cases where AI projects fail because the business context is not defined clearly. AI is fundamentally about context.

These challenges exist, but at the same time,

"AI is fundamentally about context, ensuring business context is embedded everywhere, along with unified data, is critical to creating a cohesive AI strategy."

organisations are clearly moving toward production. Within any organisation, you will still see all three phases—experimentation, proof of concepts, and production—happening simultaneously.

We recently released a study on ROI and impact, conducted with more than 2,050 respondents globally, including a large number from India. The India statistics show that 33% of respondents already have AI use cases in production. Another 33% have initiatives in motion and expect to see impact within the next 12 months. Around 24% are still in early stages, and a small percentage is lagging. The overall trend clearly shows movement toward measurable AI impact.

At the same time, this is still an evolution. Outcomes vary depending on data maturity, internal capabilities, and how deeply AI is embedded into business processes.

CIO&Leader: With agentic AI gaining momentum but ROI still under scrutiny, how do you see its adoption scaling, and what key challenges are holding enterprises back?

VIJAYANT RAI: Some of the challenges include fragmented data estates, people and process gaps, and the absence of a clear AI design in terms of business impact.

If you look at it from a high level, organisations need a more coordinated approach to AI. They need to infuse AI across all layers of the organisation instead of creating silos. They also need to ensure that business context is embedded everywhere, and that data is unified.

Different industries have different concerns. Financial services companies are focused on whether they are getting enough ROI from AI investments. Manufacturing companies are more concerned about whether they have the right skills to adopt AI.

What is common across industries is the understanding that AI is transformative if implemented correctly. At Snowflake, we try to bring predictability into the process through cost management and FinOps, helping organisations manage AI consumption and internal chargebacks.

CIO&Leader: With increasing focus on compliance, DPDP, and sovereign AI, how should enterprises approach building compliant AI architectures?

VIJAYANT RAI: This is extremely important, not just for AI but for any technology. Regulated industries must adhere to frameworks such as RBI regulations and consent regimes like DPDP.

One advantage with Snowflake is that we operate globally and have long-standing experience in markets like Europe and the US. All the controls required under GDPR are already built into the platform and are available globally.

When customers look at DPDP compliance, these controls are already in place. We also provide capabilities for data residency. In India, for example, we operate on AWS and Azure within Indian jurisdictions, ensuring compliance requirements are met.

Beyond technology, compliance also involves processes. That is where our partner ecosystem, including Deloitte, EY, and KPMG, plays a role in helping organisations implement governance frameworks effectively.

CIO&Leader: You run on hyperscalers but also compete with them. How does this relationship work?

VIJAYANT RAI: We work with all three hyperscalers globally. In India, we operate on AWS and Azure, and globally we also work with Google Cloud. We have large commitment contracts with them, and maintain strong relationships around what we host on their infrastructure. This helps ensure optimal cost dynamics for our customers.

For a Snowflake customer, the cost of compute and storage that would typically be charged separately by hyperscalers is bundled into a single pricing model. Through our partnerships, we are able to pass on both commercial and technical benefits. Given our scale as a customer for AWS, Azure, and Google, we also have a seat at the table in terms of

"Where critical thinking and strategic decisions are required, it will require human intervention, domain expertise will come from people, while AI will define how to execute."

product roadmaps and technology direction, which allows us to maintain strong alignment.

This is a key advantage. While we compete with other players in the ecosystem, each with strong technologies, the market in India still has significant room for growth. Data and AI adoption is still in its early stages for many organisations. We position Snowflake as a fully managed AI Data Cloud.

Another important advantage of operating across multiple hyperscalers is the flexibility it provides. In regulated industries, for instance, the ability to enable multi-cloud disaster recovery is highly valuable. This is possible because we operate across different cloud environments, which is not the case for many providers. It allows us to offer both flexibility and resilience, while continuing to compete where necessary.

CIO&Leader: How do partnerships with AI companies help accelerate adoption?

VIJAYANT RAI: We recently announced US $200 million partnerships with Anthropic and OpenAI. Our objective is to provide customers with access to the most suitable large language models (LLMs), aligned with their specific use cases, while ensuring a seamless consumption experience.

The platform operates on a unified consumption model using Snowflake credits, which applies consistently across both analytics and AI workloads.

These partnerships enable access to more than 15 of the world's leading LLMs. The underlying models, including those from Anthropic and OpenAI, are integrated within the platform, allowing customers to leverage them without added complexity.

Given the scale of these commercial relationships, there is also a strong engineering alignment. This ensures that advancements from these LLM providers can be efficiently integrated and delivered to customers, enabling them to benefit from continuous innovation across the ecosystem.

CIO&Leader: How do you help enterprises build responsible and trustworthy AI systems?

VIJAYANT RAI: We work closely with our partners across multiple engagements. The approach is not limited to technology or tools; it also involves process and industry-specific expertise. Partners play a critical role in bringing domain knowledge, whether in financial services, manufacturing, or other sectors. For instance, when integrating SAP data, we collaborate with partners who specialise in that ecosystem.

Several partners also work with us on the process layer, including decisions around what data is ingested and how it is managed within Snowflake. Alongside this, the platform itself includes built-in controls and technology features for security and governance.

A core principle of our platform is that it is designed to be easy, connected, and trusted, with trust at the centre. Customer data remains fully owned and controlled by the customer. We do not access or use that data; our role is to provide the framework for security and governance.

As new security capabilities or governance features are developed, they are integrated directly into the platform, allowing customers to benefit from continuous enhancements without additional overhead. This creates a combined model where technology-driven controls are embedded within the platform, while partners contribute to the process and implementation layer.

CIO&Leader: With the rise of open-source models, how should enterprises decide between sovereign stacks and global platforms?

VIJAYANT RAI: There are different considerations in each case. From the perspective of Snowflake as a global AI data cloud, a key advantage is the ability to rapidly integrate new

innovations. For example, when the open-source model DeepSeek was released, it was onboarded onto the platform within a short timeframe.

As an engineering-led organisation, the focus is on continuously integrating, testing, and operationalising new large language models. For most enterprises, building and maintaining such capabilities independently at scale is complex and resource-intensive.

Organisations that attempt to build their own infrastructure or data center environments often face challenges in keeping pace with the rapid evolution of AI technologies.

There are, however, valid use cases for sovereign or fully controlled AI stacks, particularly in sectors such as defence or highly regulated industries, where data control and governance requirements are more stringent.

The broader trend reflects increasing demand for control over data residency, security, and model deployment. In response, there has been a shift toward localised infrastructure and architectures that align with regulatory and compliance requirements, while still enabling access to evolving AI capabilities.

CIO&Leader: Where does India stand in terms of AI adoption and scalability today? What are the key strengths and challenges that will shape its growth?

VIJAYANT RAI: India stands out for its large talent pool and the scale at which digital-native companies operate. Organisations such as Swiggy serve tens of millions of users daily, demonstrating a level of scale that is rare globally.

The country has already proven its ability to build population-scale digital systems through initiatives like Aadhaar, UPI, and India Stack. AI can build on this foundation to drive large-scale impact in sectors such as financial inclusion, agriculture, and healthcare.

However, adoption is still evolving. A significant portion of enterprise data remains on-premises, particularly in regulated industries, where compliance and governance requirements slow cloud adoption. Organisations are progressing in phases, but the scale of AI deployments is increasing steadily.

Overall, the opportunity remains substantial, with strong momentum expected across sectors such as financial services, manufacturing, and digital-native businesses.

CIO&Leader: How do digital-native and traditional enterprises differ in data and AI adoption, and what distinct challenges does each face?

VIJAYANT RAI: Digital-native companies are inherently cloud-first, which means they do not face the challenge of migrating data to the cloud. As a result, they are able to adopt AI and analytics capabilities faster and at scale, leveraging platforms like Snowflake to support large user bases.

In contrast, traditional enterprises operate on legacy sysw tems with data distributed across multiple environments. Their transition to the cloud is more complex, particularly in regulated industries where compliance and governance requirements must be addressed. This results in a more gradual adoption curve.

However, traditional organisations also bring unique advantages, particularly access to long-term and diverse datasets. With the advancement of generative AI, they are now able to unlock value from unstructured data sources that were previously difficult to analyse. This is enabling new use cases, including areas like fraud detection and advanced analytics.

Ultimately, while digital-native companies move faster due to their cloud-native foundation, traditional enterprises are progressing steadily. In both cases, organisations are prioritising use cases that deliver clear business value and measurable return on investment.

CIO&Leader: How can enterprises balance data residency requirements with global LLM adoption, and how should they address evolving skill needs?

VIJAYANT RAI: From a data perspective, localisation is already addressed through in-country infrastructure. Platforms like Amazon Web Services and Microsoft Azure enable data residency within India, ensuring compliance with regulatory requirements.

On the AI side, the approach is to bring models closer to where the data resides. Many cloud environments already support AI infrastructure within India, and as demand scales, this capacity continues to expand. This allows enterprises to leverage global models while maintaining control over data.

From a workforce perspective, the industry is undergoing a generational technology shift, similar to earlier transitions such as the internet, mobile, and cloud. One key change is that AI tools have become significantly more accessible, enabling users to perform tasks that previously required specialised technical skills. This has led to measurable improvements in productivity across functions, including sales, operations, and analytics.

AI is also driving innovation by enabling organisations to process and analyse large-scale datasets that were previously unmanageable. This is opening up new possibilities across industries, including areas like research and complex data-driven decision-making.

In terms of skills, the emphasis is shifting toward critical thinking, strategic decision-making, and domain expertise. While AI can enhance execution and provide insights, human judgment remains essential in defining objectives and interpreting outcomes. Organisations should focus on enabling employees to work alongside AI systems, equipping them with the skills required to leverage these tools effectively while adapting to evolving roles. ■

Datacentres are no longer just service providers, they are custodians of critical digital ecosystems

Anil Nama, CIO at CtrlS Datacenters outlines the shift to resilience-driven data centres, using AI-led observability, zero-trust security, and sovereign compliance to ensure uninterrupted operations

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

As enterprises accelerate their shift toward AI-led, distributed digital infrastructure, the role of data centres is undergoing a fundamental transformation. No longer confined to being secure, isolated facilities, modern data centres now underpin critical national workloads, from financial systems to digital public platforms, making them central to both business continuity and national resilience.

In this conversation, **Anil Nama, CIO at CtrlS Datacenters** outlines how CtrlS is re-architecting its data centre framework to align with these new realities. The discussion spans its transition to resilience-driven, zero-trust infrastructure; the integration of AI-led observability and automated response systems; and the complexities introduced by high-density AI workloads. It also addresses the growing importance of sovereign security, hardware supply chain integrity, and the convergence of physical and cyber security.

CIO&Leader: Traditionally, data centre security focused on "keeping the bad actors out." With the rise of distributed workloads and AI, the focus has shifted to resilience, how quickly you can recover. How is CtrlS re-architecting its framework to ensure that uninterrupted service is guaranteed even during an active breach attempt?

ANIL NAMA: The shift from perimetre-based security to resilience-driven architecture reflects the evolving complexity of modern digital infrastructure. At CtrlS, the focus is no longer just on preventing breaches, but on ensuring uninterrupted service even in the event of an active threat.

Our Rated-4 fault-tolerant infrastructure eliminates single points of failure, ensuring continuous operation during disruptions or cyber incidents. It is built to contain, isolate, and recover from threats without affecting service. We use real-time monitoring and AI-led observability to detect anomalies early and trigger automated responses like workload isolation, traffic rerouting, and rapid recovery failover.

Our framework is built on zero-trust principles, continuous validation, and automated failover mechanisms that ensures any breach is contained, minimising its impact across environments. By integrating disaster recovery, active-active deployments, and intelligent traffic rerouting, we ensure that workloads remain operational even during disruptions. Cyber resilience, for us, is about designing systems that can absorb disruptions, respond in real-time, recover instantly, and continue delivering seamless performance without impacting end-users.

CIO&Leader: As data centres host more national digital workloads, they effectively become part of a country's critical infrastructure. How does this change the compliance and sovereign security obligations for an operator like CtrlS compared to five years ago?

"AI is fundamentally about context, ensuring business context is embedded everywhere, along with unified data, is critical to creating a cohesive AI strategy."

ANIL NAMA: As datacentres increasingly support critical national infrastructure that ranges from financial systems to digital public platforms, the expectations around compliance and sovereign security have significantly intensified. Compared to five years ago, operators are now required to align more closely with evolving data localisation norms and sector-specific regulations, but also with stringent audit frameworks and real-time compliance validation.

At CtrlS, this meant strengthening our compliance posture across multiple layers, including globally recognised standards such as ISO 27001, ISO 22301, and ISO 27701, supported by rigorous operational processes and regular internal and external audits.

At the same time, sovereign security now plays a crucial role in the design and operation of datacentres. This involves keeping data within national borders, maintaining complete control and oversight of data environments, and ensuring that infrastructure aligns with changing national data governance policies. Datacentres are no longer just service providers, they are custodians of critical digital ecosystems where trust, security, and accountability are paramount.

CIO&Leader: AI requires high-density compute and massive data ingress/egress. How has the specialised infrastructure required for AI (GPUs, liquid cooling, high-speed interconnects) created new, unforeseen vulnerabilities in the physical or digital data centre layer?

ANIL NAMA: AI infrastructure introduces a new class of challenges that extend beyond traditional datacentre risks. High-density GPU environments, liquid cooling systems, and high-speed interconnects significantly increase both the physical and digital complexity of operations, creating potential vulnerabilities.

For instance, dense GPU clusters can create thermal hotspots undetected by standard monitoring, and reliance on advanced cooling makes system vulnerable to operational or malicious disruptions stability. At the same time, liquid cooling presents additional physical risks, including system dependencies, leak risks, and stricter infrastructure tolerances design.

On the digital side, large-scale data movement and multi-tenant AI workloads heighten risks around data integrity, model security, and lateral threat propagation.

CtrlS addresses these challenges by combining secure infrastructure design with advanced monitoring, micro-segmentation, and real-time threat detection. The focus is on ensuring that as infrastructure becomes more powerful and dense, it also becomes more controlled, observable, and resilient against emerging threat vectors.

CIO&Leader: Is CtrlS currently deploying "AI-to-fight-AI"? Specifically, how are you using machine learning to detect anomalous patterns in data traffic or power consumption that might signal a sophisticated, low-and-slow cyber-attack?

ANIL NAMA: Yes, AI is increasingly becoming a critical component of our cybersecurity strategy. At CtrlS, machine learning models are used to analyse large volumes of telemetry data across network traffic, system performance, and even power consumption patterns to detect anomalies in real time. This multi-layer visibility enables us to detect subtle anomalies that could suggest early-stage or evasive attack patterns.

This approach is unique because it can correlate signals across various infrastructure layers. For example, analysing deviations in power consumption or workload together with network patterns can enable the identification of subtle, low-and-slow attacks that may otherwise go unnoticed in traditional systems. By correlating behavioural patterns across infrastructure layers. AI-driven systems can then flag deviations, trigger automated responses such as traffic isolation, workload containment, and rapid remediation, significantly reducing response time and limiting potential impact.

As cyber threats become more sophisticated and AI-powered, leveraging AI for defence is essential to

maintain speed, accuracy, and scale in threat detection and response.

CIO&Leader: In a Rated-4 environment, we often talk about redundant power and cooling. However, how do you integrate physical security (biometrics, surveillance, access control) with cyber security (zero-trust network access) so they aren't operating in silos?

ANIL NAMA: In modern datacentres, physical and cyber security can no longer operate in silos, they must function as a unified, intelligence-driven system. At CtrlS, we follow a converged security model where physical access controls, surveillance systems, and biometric authentication are tightly integrated with cyber security frameworks such as zero-trust access and continuous monitoring.

This integration is enabled by a multi-layered, zonal security architecture aligned with our Rated-4 design, which connects each physical access point to a digital identity and system-level activity. For instance, physical access events are correlated with system-level activity in real time across user authentication, workload access, and network behaviour, enabling us to detect anomalies such as mismatches between physical presence and system activity. Multi-layered security zones, 24/7 security operations centres, and integrated monitoring platforms ensure complete visibility across both physical and digital environments. This unified approach strengthens overall resilience by enabling faster detection, coordinated response, and comprehensive risk management across the entire datacentre ecosystem.

CIO&Leader: With global geopolitical tensions, the provenance of hardware from chips to cooling units, is under scrutiny. How does CtrlS ensure the integrity of its hardware supply chain to prevent hardware-level backdoors in the data centre?

ANIL NAMA: Ensuring hardware integrity has become a critical priority in today's geopolitical environment. At CtrlS, we address this through a stringent vendor selection process, sourcing from trusted global OEMs and certified suppliers with proven track records.

We implement rigorous validation, testing, and auditing mechanisms across the supply chain to verify hardware authenticity and performance before deployment. Additionally, secure procurement practices, lifecycle management, and controlled access to critical infrastructure components help mitigate risks associated with hardware-level vulnerabilities. The focus is on building a transparent, traceable, and secure supply chain that minimises exposure to potential backdoors or compromised components.

> "One of the most common misconceptions is that moving to a Rated-4 datacentre transfers full security responsibility to the provider. Security is a shared responsibility. "
>
> **Anil Nama**
> CIO, CtrlS Datacenters

CIO&Leader: Many enterprises assume that moving to a Tier-4/Rated-4 data centre absolves them of certain security burdens. Where do you see the most common security gaps in the hand-off between the data centre operator and the enterprise tenant?

ANIL NAMA: One of the most common misconceptions is that moving to a Rated-4 datacentre transfers full security responsibility to the provider. Security is a shared responsibility. While CtrlS ensures infrastructure-level security, resilience, and compliance, enterprises remain responsible for securing their applications, data, and access controls.

Key gaps often arise in areas such as identity management, misconfigurations, inadequate encryption, and a lack of visibility across hybrid environments. To address this, we work closely with customers to provide clear demarcation of responsibilities, best-practice guidance, and integrated security frameworks, supported by continuous monitoring and managed services. Bridging this gap requires collaboration, continuous monitoring, and alignment between infrastructure and application-level security strategies.

CIO&Leader: Looking toward 2027-2030, what is the one emerging threat that isn't currently on the average CIO's radar, but is keeping data centre operators like yourself awake at night?

ANIL NAMA: One emerging threat that needs greater attention is the rise of AI-driven, autonomous cyberattacks targeting infrastructure at scale. These attacks could be capable of adapting in real-time, exploiting vulnerabilities across interconnected systems, and operating at speeds beyond human response.

At the same time, risks around data integrity, such as model poisoning and manipulation of training datasets, could have far-reaching consequences, especially as AI becomes embedded in critical decision-making systems. In highly interconnected datacentre environments, such threats could also cause cascading effects throughout compute, network, and infrastructure layers.

For datacentre operators, this means preparing for a future where threats are not only more sophisticated but also more systemic. Building intelligent, adaptive, and self-healing infrastructure will be key to staying ahead of this evolving risk landscape. ■

The journey is moving towards a self optimising enterprise

Ritwik Batabyal, Chief Technology and Innovation Officer at Mastek highlights how enterprises can move beyond AI pilots by reengineering processes, funding innovation through efficiency gains, and building self-optimising, AI-led organisations.

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

As enterprises move beyond AI experimentation, the focus is shifting toward scalable, outcome-driven transformation. Organisations are under increasing pressure to convert AI investments into tangible business value, while navigating evolving paradigms such as agentic AI, data sovereignty, and intelligent automation at scale.

In this conversation, **Ritwik Batabyal, CTO and Innovation Officer at Mastek,** explains how the company's "Lead with AI" approach is enabling this shift by reengineering core business processes, embedding intelligence, and validating solutions internally before scaling. He also shares insights on legacy modernisation, monetising AI through efficiency gains, India's deep-tech opportunity, and the rise of self-optimising enterprises—what he describes as the next phase of enterprise evolution.

CIO&Leader: Mastek has recently rebranded its approach as "Lead with AI." How are you moving beyond the pilot stage, which many enterprises struggle with, and delivering the 20–30x ROI you claim from AI-led programs?

RITWIK BATABYAL: Before discussing outcomes, it is important to explain the process.

We begin by selecting specific internal business processes such as people supply chain, accounts receivable, or accounts payable. These are broken down into journey touchpoints using our ADOPT framework. Each touchpoint is then reimagined using AI or intelligent automation.

We transform the entire process against defined KPIs to achieve measurable outcomes. This transformation is first implemented internally to validate its effectiveness. Once proven, we offer it externally as a software-driven service.

This internal validation allows us to deliver consistent outcomes in external deployments. Therefore, the 20–30x ROI is not a promise but a result of a repeatable engineering discipline. It is essentially a "pilot in action" approach rooted in reengineering business processes.

CIO&Leader: You have secured significant contracts with the UK Home Office and other public sector entities. In such high-stakes environments, how does your approach ensure that legacy modernisation does not simply migrate inefficiencies to the cloud but actually transforms service delivery?

RITWIK BATABYAL: Most organisations approach legacy modernisation as a migration exercise. We take a different approach by reimag-

ining citizen journeys rather than focusing solely on system transformation.

We simplify decision layers, introduce intelligent feedback loops at critical points, and automate compliance-heavy processes. This shifts modernisation from a technical migration exercise to a measurable improvement in speed, transparency, and trust.

Our focus is on outcomes driven by user experience, particularly from a citizen-centric perspective.

CIO&Leader: The industry is moving from chatbots to agentic AI systems capable of executing multi-step workflows. What architectural challenges do enterprises face in scaling such systems, and how can they manage autonomous decision-making while maintaining trust?

RITWIK BATABYAL: Agentic AI is not a standalone technology; it is a combination of multiple interconnected elements. It requires orchestration and the ability to build trust.

Many enterprises struggle because they treat it as a product deployment rather than an integrated system design. Our approach focuses on identifying both data in motion and data at rest, defining implicit decision rules, and embedding governance into the system.

We incorporate observability to monitor agent behavior and decisions, along with human-in-the-loop mechanisms where necessary. These elements collectively ensure orchestration, accountability, and trust.

"The 20–30x ROI is not a promise but a result of a repeatable engineering discipline."

CIO&Leader: Cost and revenue remain central concerns. Many enterprises are unable to move beyond experimentation to generate meaningful returns from AI. What is the most effective monetisation strategy for mid-market enterprises?

RITWIK BATABYAL: The strategy must be sequential.

The first step is internal cost optimisation. By improving efficiency in business processes, organisations can generate measurable ROI. These gains can then fund innovation.

The next step is reinvesting those savings into AI-driven products or new revenue streams. In essence, innovation should be funded through efficiency gains. This approach ensures sustainability and reduces dependency on external funding.

CIO&Leader: With increasing regulatory pressure around data localisation and sovereignty, how can enterprises balance centralised data platforms with localised compliance requirements?

RITWIK BATABYAL: The future lies in building systems that are globally intelligent but locally compliant.

Data management should leverage centralised intelligence while maintaining strict localised control. This can be achieved through policy-driven data layers, strong encryption, and region-specific governance models.

Such an approach ensures that data remains usable at a global level while adhering to local regulatory requirements.

CIO&Leader: You have experience across startups and large enterprises. What is the key to building a scalable innovation ecosystem, and how do you select partners?

RITWIK BATABYAL: We evaluate partners through two primary lenses.

The first is readiness, which includes:

- Enterprise alignment with the use case
- Depth of domain expertise
- Agility in integration

The second is the operational model, which includes:

- Willingness to invest in early-stage innovation (MVP development)
- Transparency in creating a mutually beneficial commercialisation model

Innovation cannot be purely commercial from day one. Both parties must invest initially to build a sustainable, scalable solution.

CIO&Leader: India aspires to become a deep-tech hub. What is the one critical gap that needs to be addressed?

RITWIK BATABYAL: India's strengths lie in two areas: service orientation and human capital.

These strengths must be converted into monetisable intellectual property. Historically, the IT industry has focused on labor arbitrage, but the future lies in building IP-driven solutions.

Transforming service capabilities and talent into scalable, commercially viable IP is essential for India to lead in deep tech.

CIO&Leader: Many industry leaders believe the cloud era is maturing and the focus is shifting toward intelligence. Is there a third transformational phase emerging?

RITWIK BATABYAL: The trajectory is moving toward what I would describe as a concept of singularity.

If you look at the evolution, we moved from cloud transformation to broader digital transformation, and now into AI-led transformation. However, beyond these phases, the direction is increasingly toward a self-optimising enterprise.

AI, in this context, is an enabler rather than the end state. The real shift is toward organisations that can continuously optimise themselves. If you reflect on my earlier responses, many of them are centered on optimising internal processes to unlock efficiency, which then enables the creation of new revenue streams.

This continuous internal optimisation is what will define the next phase. Enterprises that can adopt this model effectively, while ensuring that appropriate guardrails, governance, and controls are in place—will emerge as sustainable winners in the long term.

CIO&Leader: You mentioned the concept of singularity. Could you elaborate on what exactly you mean by that for better clarity?

RITWIK BATABYAL: When I refer to singularity in this context, I am essentially talking about a self-sustaining and self-optimising enterprise.

Traditionally, innovation and transformation have relied on a triad involving academia, industry, and customer use cases. There has always been a collaborative architecture where different entities contribute to building and refining solutions.

"Innovation should be funded through efficiency gains, ensuring sustainability and reducing dependency on external investment."

However, what is now emerging is a shift at the level of business processes themselves. Organisations are beginning to look inward, identifying their internal processes and systematically optimising them in a stepwise manner.

As these internal processes become more intelligent and efficient, the enterprise starts evolving into a more self-reliant entity. This reduces the dependency on external collaborative constructs and enables the organisation to function as a more unified and self-driven system.

That is what I mean by singularity—an enterprise that is capable of driving its own optimisation, learning, and evolution from within, rather than relying predominantly on external ecosystems. ■

AI agents will become core, but governance will decide who scales

Ajay Ajmera, CIO of Rockman Industries, advocates a governance-first approach to AI adoption, where agents are deployed in controlled pilots across key functions such as IT, finance, procurement, cybersecurity, and manufacturing

By **Musharrat Shahin** | editor.tech@timesgroup.com

Rockman Industries, a leading Indian auto-components manufacturer and part of the Hero Group, is taking a measured, governance-first approach to AI adoption. Ajay Ajmera, CIO of the company, explains how AI agents are being piloted across IT, finance, procurement, cybersecurity, and manufacturing, with a strong focus on trust, explainability, and operational risk.

Unlike organisations rushing toward full autonomy, Rockman prioritises secure integration, clearly defined risk boundaries, and embedding accountability directly into AI systems. Ajmera outlines how AI is evolving from an experimental tool to a disciplined digital execution layer that complements human judgment and strategic oversight.

CIO&Leader: Where are you today on AI agents? Are you still experimenting, running pilots, or operating AI agents in live production? Which business functions are using them today?

AJAY AJMERA: We are currently operating in a structured pilot phase. While the broader industry conversation has sprinted toward full autonomy, we've intentionally chosen a more measured approach. Today, our AI agents are in controlled pilots across critical functions: IT service management, finance analytics, procurement insights, and cybersecurity monitoring. We're also testing predictive signals and quality pattern detection in our manufacturing operations.

Right now, it's not about handing over decision-making authority. It's about evaluating behavior, accuracy, and integration complexity and ultimately, earning the trust of our teams. We view these pilots not merely as technical proofs of concept, but as governance rehearsals. We are stress-testing our control frameworks and risk thresholds just as rigorously as the AI itself. We're piloting with clear intent to achieve secure, sustainable production at scale.

CIO&Leader: Where do you draw the line on autonomy? Which decisions can AI agents make on their own today, and which decisions are strictly off-limits for machines?

AJAY AJMERA: I look at AI autonomy as a risk and accountability issue, not just a technological capability. Our framework for delegating decisions comes down to impact and reversibility. Today, agents act independently only in structured environments where their actions are measurable and correctable, such as triggering predefined workflows, prioritising cases, or executing small-value transactions.

However, decisions that carry material financial, regulatory, or human consequences remain strictly off-limits. Hiring, terminations, strategic investments, and vendor negotiations require nuanced judgment that machines simply don't possess. Our operating principle is straightforward: if an error carries high or irreversible consequences, a human must retain final accountability. AI manages scale and volume beautifully, but true judgment and ethi-

"Right now, it's not about handing over decision-making authority. It's about evaluating behavior, accuracy, and integration complexity—and earning the trust of our teams."

cal context remain uniquely human. Autonomy will expand in our organisation, but only as explainability and governance mature alongside it.

CIO&Leader: What level of risk have you delegated to AI? In practical terms, what is the maximum financial, operational, or reputational impact you are comfortable allowing an AI agent to carry?

AJAY AJMERA: We take a highly structured approach to risk delegation. At our current maturity level, we cap AI autonomy at low-to-moderate risk within tightly governed, reversible parameters. From a financial standpoint, AI operates strictly within predefined transaction thresholds; any material exposure automatically defaults to human authorisation.

Operationally, we leverage AI to optimise workflows and prioritise cases, but we firmly restrict it from actions that could disrupt core production or alter strategic commitments. We are most conservative regarding reputational risk; public-facing decisions and regulatory submissions remain entirely human-driven. Our guiding heuristic is straightforward: if an error cannot be rapidly reversed—financially, operationally, or reputationally—it is not delegated to a machine. While these boundaries will naturally expand as our governance frameworks mature, today our delegation is cautious and continuously monitored. Ultimately, autonomy without a defined risk appetite isn't innovation; it's exposure.

CIO&Leader: If asked, can you trace and explain how autonomous decisions were made and who approved the rules behind them? Could you explain your AI decisions to a regulator tomorrow?

AJAY AJMERA: Our position is uncompromising: if an AI's decisions cannot be explained, it does not get deployed. Even in our current pilot phase, every agent operates under strict version control, documented rules, and clear business ownership. We maintain comprehensive decision logs, ensuring we can trace not only the output of the AI, but the underlying policy and the specific human authorisation behind it. Ultimately, we treat AI decisions with the exact same level of auditability as our financial decisions.

If a regulator walks in tomorrow, we are fully prepared to articulate the agent's purpose, its data supply chain, the governing decision logic, and our oversight mechanisms. While no system is flawless, explainability is a foundational design principle for us. The fundamental shift for modern CIOs is realising that governance must be embedded directly into the AI lifecycle from day one. Transparency is no longer optional; it is the fundamental license to operate.

CIO&Leader: As AI agents connect to ERP, CRM, core banking, or manufacturing systems, what safeguards ensure control, visibility, and auditability? How do you integrate agents into core systems without losing control?

AJAY AJMERA: The moment an AI agent connects to platforms like ERP or core banking, it ceases to be an experiment and becomes a foundational part of your operational backbone. For us, secure integration is rooted in architectural control. AI agents never get unrestricted access.

We essentially manage them as digital employees operating under robust governance layers. This means enforcing strict role-based access controls (RBAC) and ensuring they interact solely via controlled APIs. We implement hard financial caps, mandate segregation of duties, and maintain real-time monitoring dashboards with immediate override capabilities. In manufacturing, agents might trigger alerts or suggest optimisations, but they cannot independently alter production parameters beyond tightly approved tolerance bands. Every action is fully traceable—we know who configured it, what rule it followed, and what data it used. You don't lose control by

integrating AI into core systems; you only lose control when governance is treated as an afterthought.

CIO&Leader: What matters most in your model today: monitoring, human override, policy rules, or strict limits on what agents can do? How are you redefining 'control' as machines start acting?

AJAY AJMERA: While all four elements are essential, I prioritise policy boundaries and strict operational limits above all else. Overrides and monitoring are necessary, but they are fundamentally reactive. True security lies in proactive design—architecting the system so the agent simply cannot operate outside clearly defined guardrails. We establish those absolute boundaries first, and only then do we implement continuous monitoring and real-time override capabilities.

We are fundamentally redefining 'control' as machines begin to act. Historically, control meant manual approvals at every procedural choke point. Today, control is architectural. It is embedded in the system through segregation of duties, policy-driven automation, and transaction caps. While the human 'kill-switch' remains vital, our objective is to elevate supervision from an operational task to a strategic one. Control is shifting from being procedural to being systemic. That is the true evolution for enterprise leadership today: transitioning from managing processes to engineering accountability directly into the system.

CIO&Leader: Which decisions do you realistically see machines taking over in your industry, and which will always require human judgment? What will AI agents be trusted to decide in three years?

AJAY AJMERA: Looking ahead three years, I anticipate AI agents confidently owning data-heavy, pattern-driven decisions. In manufacturing and enterprise operations, this translates to autonomous inventory optimisation, predictive maintenance scheduling, fraud pattern detection, and dynamic pricing within strict thresholds. These environments are rich in measurable variables and perfectly suited for continuous algorithmic optimisation.

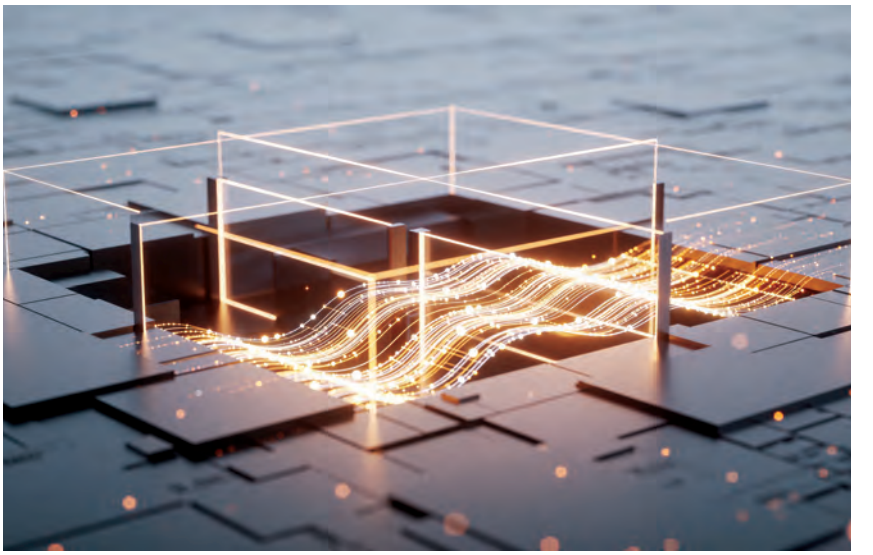

"AI will function as a digital execution layer, trusted to manage massive operational complexity at scale while humans manage direction and values."

However, decisions requiring contextual nuance, ethical weighting, and long-term strategic foresight will remain strictly human-led. Machines simply lack the situational awareness and accountability required for capital allocation, regulatory interpretation, talent management, or crisis response. In the near future, AI will effectively function as a digital execution layer, trusted to manage massive operational complexity at scale while humans manage direction and values.

CIO&Leader: Will AI agents become a core execution layer in your enterprise, or will scale stall due to regulation, architecture limits, or cost pressures? Five years out, what's your honest view?

AJAY AJMERA: If I'm being completely honest, AI agents will inevitably become a core execution layer—but through a measured, incremental evolution rather than an overnight revolution. The technology itself will not be the limiting factor. The true constraints will be regulatory frameworks, data architecture readiness, the cost of large-scale inference, and internal trust. Regulation will rightfully slow reckless adoption, which is a net positive for the industry. Furthermore, legacy systems that aren't API-ready will fundamentally struggle to integrate these autonomous layers.

Looking five years out, I don't foresee scale stalling; rather, I see it maturing into highly disciplined adoption. AI will orchestrate a massive portion of operational execution, freeing human leaders to focus entirely on strategy and high-stakes decisions. The ultimate winners won't be measured by the sheer volume of deployed agents. They will be the organisations that embedded observability and accountability into their foundations from day one. AI agents will become core to the enterprise, but governance will dictate who successfully scales and who stops. ■

ET Edge | CISOFORUM

PRESENTS

CISO FORUM CONFERENCE '26 MID YEAR

Building A Cyber Resilient Enterprise

12 -14 June 2026 ▪ Hayatt Regency Jaipur Mansarovar

Where India's RS. 25,000 Cr Cybersecurity Spend Gets Decided.
Be in the room. Or read about it later.
3-day Residential Experience with 100+ decision-makers.

Who's in the Room

▪ 100+ CISOs & Security Heads ▪ 42% BFSI & Financial Services ▪ 24% IT / ITeS & SaaS
18% Manufacturing & Pharma ▪ 16% Retail, Telecom & Conglomerates
▪ Average security spend 50 Cr (approx)

What You Walk Away With

▪ 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

Fireside chat with with Legendary
Krishnamachari Srikkanth

Early-bird offer closing soon. few slots available

For Sponsorship, Write to:

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

Sourabh Dixit
Chief Manager - Projects,
sourabh.dixit@timesgroup.com,
+91 9971475342

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

Subhadeep Sen
Sr. Project Manager,
subhadeep.sen@timesgroup.com,
+91 9611307365

SILVER PARTNERS: netskope, rubrik

EXHIBIT PARTNER: Barracuda. Your business, secured.

CONCEPT BY: ET Edge | CISOFORUM

RNI No. DELENG/2012/42736

AMD PRESENTS

27th Annual Conference
CIO&LEADER
The Agentic Enterprise

PLATFORMS. PEOPLE. POLICY. PROFITS.

30th July - 2nd August 2026 • Jaipur

200+
CIOs & future CIOs, managing IT budgets between ₹50 Cr - ₹2,000 Cr annually.

85%
from ₹5,000 Cr+ enterprises.

₹4.2
lakh crore in enterprise IT spend.

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

Legacy: CIO&Leader Annual Conference

26 EDITIONS

16 YEARS OF NEXT100

100+ PARTNER BRANDS

Early-bird offer closing soon. Few slots available.

PRESENTING PARTNER AMD

CONCEPT BY ET Edge | CIO&LEADER | INDIA'S ELITE FUTURE CIOs NEXT100

For Sponsorship, Write to

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

Sourabh Dixit
Chief Manager - Projects
sourabh.dixit@timesgroup.com
+91 9971475342

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

Subhadeep Sen
Sr. Project Manager
subhadeep.sen@timesgroup.com
+91 9611307365

#CIOandLeaderConference