

CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

JUNE 2026 • VOLUME 15 • ISSUE 03 • ₹150



THE NEW CIO

How AI agents, autonomous workflows and intelligent decision-making are reshaping the role of the CIO. Pg. 10

A 9.9 GROUP
PUBLICATION

cioandleader.com
[in](#) [cioandleader/](#)
[f](#) [cioandleader](#)

Top Row (L-R): **AJ Sunder**, COP/CIO & CO-founder, Responsive; **Vinod Krishnan**, CIO, Birlasoft;
Anjaiah Surgi, CIO, Hindware; **Chitti Babu**, Group CIO, Aurobindo Pharma, **Nikhil Prabhakar**, CIO, IndiaMart
Bottom Row (L-R): **Rakesh Bhardwaj**, Global CIO, Lupin; **Rajeev Soota**, Vice President-IT, Usha International;
Ashok Jade, Global CIO, Kirloskar Brothers



ET
studio**talks**



ET
studio**talks**



ET EDGE | CIO&LEADER STUDIOTALKS WHERE TECHNOLOGY MEETS THE SPOTLIGHT !

ET Edge | CIO&Leader 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 top 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 ET Edge | CIO&Leader StudioTalks.

SECURE YOUR SPOT NOW !

For more information
Jatinder Singh
Chief Editor, Enterprise Tech Publications
ET Edge - The Times Group
Jatinder.singh@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



Are you ready to lead a workforce that doesn't exist today?

Imagine walking into the boardroom in 2030. Your enterprise hasn't just hired more people; it has onboarded thousands of AI agents. They work across finance, HR, supply chain, engineering, cybersecurity, customer service, and every other business function. They don't simply answer questions or draft emails. They execute workflows, negotiate with other agents, access enterprise systems, and make decisions within clearly defined guardrails.

For years, CIOs have been measured by how well they deployed technology. That era is giving way to a far more complex mandate: governing a digital workforce.

In conversations with CIOs over the past few months, one theme has surfaced repeatedly. The real challenge isn't building AI agents; it is managing them. Who owns an AI agent? Who grants and revokes its access to enterprise data? Who monitors its actions, audits its decisions, and takes accountability when something goes wrong? How do you prevent thousands of autonomous agents from becoming the next generation of insider risk?

These are not just technology questions. They are questions of governance, leadership, operating models, and trust.

The conversation becomes even more interesting when we ask a more fundamental question: What exactly is an AI agent? Is it a digital co-worker capable of reasoning, planning, and autonomous action? Or is it, despite its growing sophistication, still a predictive model that recognizes patterns and follows instructions? The answer matters because enterprises are beginning to delegate work, not just tasks, to these systems.

History tells us that technology alone hardly transforms organizations. Cloud didn't. Digital transformation didn't. Data platforms didn't. Enterprises changed only when people, processes, culture, and governance evolved alongside the technology. Agentic AI will be no different.

The biggest challenge isn't whether AI agents are ready for the enterprise. It's whether enterprises are ready for AI agents. Are our policies designed for autonomous decision-makers? Can our security architectures accommodate thousands of non-human identities? Are our risk frameworks prepared for AI-to-AI interactions? And are leaders ready to manage teams where most workers may not be human at all? ■

“The biggest challenge isn't whether AI agents are ready for the enterprise. It's whether enterprises are ready for AI agents.”



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



COVER STORY

10-19

The New CIO

How AI agents, autonomous workflows and intelligent decision-making are reshaping the role of the CIO.



Cover Design by:
Shokeen Saifi



Please Recycle this
Magazine and Remove
Inserts before Recycling

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



NEWS & VIEWS

08

Why Google Antigravity 2.0 signals towards a multi-agent...



SPECIAL REPORT

20-23

COMPUTEX 2026: From Personal Devices to AI Infrastructure



INSIGHT

24-26

How Synology is reinventing itself for the AI era



27-28

The CEO agenda for 2026: AI, agility and workforce transformation



TECH TALK

29-31

Any AI strategy must be tightly connected to a...

BY ROBERT PIZZARI



32-34

Many companies are still building human-optimised applications

BY MIKE ANDERSON



35-37

AI without guardrails is just risk at scale

BY ROHIT SHUKLA



38-40

IT has moved beyond keeping the lights on!

BY ANJIAH SURGI

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**



CIO MOVEMENTS



Pradipta Patro Appointed CISO & VP – Data Privacy at Polycab India

Pradipta Patro joins Polycab India from RPG Group to lead cybersecurity strategy, strengthen data privacy frameworks, and enhance enterprise resilience and governance capabilities.



Srijesh Kumar Appointed Global CPTO at Practo

Srijesh Kumar joins Practo from Salesforce as Global Chief Product and Technology Officer to drive healthcare innovation, scalable platforms, and digital transformation initiatives.



Atul Ahuja Appointed Area VP & GM – India at Elastic

Atul Ahuja joins Elastic from P E E R A Information Technology Consultants to lead India operations, enterprise growth, and AI-driven market expansion initiatives.



Mohit Kalra Elevated as VP-2 & CISO at ORIX India

Mohit Kalra is promoted at ORIX India to strengthen cybersecurity governance, enterprise resilience, compliance frameworks, and organization-wide risk management initiatives.



Simon Davies Appointed SVP & GM – APJ at Databricks

Simon Davies joins Databricks from SAP to lead Asia Pacific and Japan growth, accelerate enterprise AI adoption, and expand regional business operations.



Jayant Magar Appointed CIO at Tata Advanced Systems

Jayant Magar joins Tata Advanced Systems from Hindustan Platinum to lead enterprise IT strategy, cybersecurity, and digital transformation initiatives across aerospace and defence operations.



Ashish Faujdar Appointed Global CISO at Lupin

Ashish Faujdar joins Lupin from Mahindra Group to strengthen global cybersecurity strategy, enterprise resilience, data protection frameworks, and secure digital transformation initiatives.



Jimit Shah Appointed CISO at LIC Housing Finance

Jimit Shah joins LIC Housing Finance from Aditya Birla Housing Finance to lead cybersecurity strategy, governance frameworks, enterprise risk management, and digital trust initiatives.



Smitesh Valanju
Appointed Global
CISO at Tata
Communications

Smitesh Valanju joins Tata Communications from Tata Digital to lead global cybersecurity strategy, enterprise resilience, digital trust, and secure infrastructure transformation initiatives.



Kalyan Achyutuni
Appointed CIO at Robi
Axiata PLC

Kalyan Achyutuni joins Robi Axiata from Nokia to drive AI-led transformation, enterprise IT modernization, cybersecurity strategy, and next-generation telecom technology initiatives.



Rishi Sareen Appointed
COO at DTDC Express

Rishi Sareen transitions from CIO to Chief Operations Officer at DTDC Express to lead enterprise transformation, operational excellence, and scalable supply chain growth initiatives.



NEXT100 Winner Hemang Doshi
Appointed Director
– IT Infrastructure &
Security at Hitech i2i

Hemang Doshi takes on an expanded leadership role to strengthen cybersecurity, compliance, resilient infrastructure, and enterprise technology transformation initiatives globally.



Ajitsingh Nawale
Appointed VP – IT at
Sterling Green Power
Solutions

Ajitsingh Nawale joins Sterling Green Power Solutions from Advik Hi-Tech to lead digital modernization, cybersecurity frameworks, and scalable enterprise technology transformation initiatives.



Parminder Singh
Appointed Founding
CEO at Reliance
Enterprise Intelligence

Parminder Singh joins Reliance Enterprise Intelligence from ClayboxAI to build AI-led enterprise platforms, drive innovation, and accelerate next-generation digital transformation initiatives globally.



Claudius Crasto
Elevated as EVP –
Innovations, Digital,
Design & AI at Bajaj Life

Claudius Crasto is elevated at Bajaj Life to drive enterprise innovation, AI-led transformation, digital ecosystems, and customer experience modernization initiatives.



Shivank Bansal
Appointed Deputy CIO
at BSEIndia

Shivank Bansal joins BSEIndia from IDFC FIRST Bank to drive AI-led infrastructure, cloud-native transformation, and scalable digital market ecosystem initiatives.

Anthropic leases Musk's Colossus Data Center for over US \$40 Billion

The agreement gives Anthropic immediate access to one of the world's largest AI computing clusters as demand for its Claude models continues to surge.

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

A securities filing has revealed one of the largest infrastructure agreements in the history of artificial intelligence. Anthropic has agreed to pay xAI, the artificial intelligence firm founded by Elon Musk, US \$1.25 billion per month to lease computing infrastructure. The arrangement runs through May 2029, bringing the total value of the contract to more than US \$40 billion.

Details of the multi-billion-dollar deal emerged not from the companies themselves, but within the initial public offering (IPO) prospectus filed with the U.S. Securities and Exchange Commission by SpaceX. The aerospace manufacturer acquired xAI in an all-stock transaction, absorbing the unit into its new SpaceX AI division.

The contract grants Anthropic full access to

the 300-megawatt capacity of the Colossus 1 data center located near Memphis, Tennessee. The facility houses more than 220,000 Nvidia graphics processing units (GPUs). The agreement also guarantees Anthropic capacity rights in the upcoming Colossus 2 expansion, which plans to scale up Nvidia GB200 hardware. To ease the technical transition, Anthropic will receive a discounted rate during the first two months while the data center completes its infrastructure ramp-up.

The financial disclosures within the filing highlight why both competitors agreed to the partnership. xAI experienced significant financial strain last year, reporting an operating loss of US \$6.4 billion on US \$3.2 billion in revenue. At the same time, user engagement for its flagship AI assistant, Grok, declined. This drop left xAI with expensive, idle hardware. By renting the cluster to Anthropic, xAI converts its excess capacity into immediate recurring revenue ahead of the SpaceX public market debut.

Anthropic faced the opposite problem. Surging demand for its Claude platform, particularly the premium Claude Pro, Claude Max, and Claude Code tools, triggered capacity shortages and slowed user service during peak hours. Rather than spending years and billions of dollars building physical data centers from scratch, Anthropic chose to lease immediate power from a direct competitor.

The contract contains a flexible exit clause that allows either party to terminate the agreement with 90 days' notice. If the partnership holds, the companies have discussed longer-term projects, including the development of orbital AI computing satellites by 2028 to bypass ground-based electrical power and cooling constraints. ■



By renting the cluster to Anthropic, xAI converts its excess capacity into immediate recurring revenue ahead of the SpaceX public market debut.

Microsoft and EY invest US \$1Bn to move AI beyond corporate pilots

Microsoft and EY are betting US\$1 billion that the next phase of enterprise AI won't be about pilots, it will be about large-scale operational transformation.

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

Microsoft and EY have expanded their long-standing alliance through a joint investment of more than US \$1 billion over the next five years. The initiative aims to help businesses move artificial intelligence past the experimental phase and deploy it across core corporate operations.

The strategy combines Microsoft's engineering teams with EY industry consultants. These uni-

fied teams will build and deploy industry-specific AI solutions tailored to individual business needs. The project uses Microsoft's specialized engineering approach to accelerate adoption across change management delivery models. Initial roll-outs will target finance, tax, risk, human resources, and supply chain management within sectors like energy, financial services, consumer goods, government, and healthcare.

EY serves as the primary testing ground, or "Client Zero," for the technology. The firm initially deployed Microsoft Copilot to 150,000 users, which yielded a 15% increase in productivity. Following these internal tests, EY is now scaling Copilot to its entire global workforce of more than 400,000 people via Microsoft 365 E7: The Frontier Suite.

Internal data from EY shows measurable operational shifts from its earlier AI deployments:

- Finance operations: Using the Microsoft Power Platform and intelligent agents reduced operational costs by 37% and cut lead times by 95%.
- Audit and assurance: EY embedded a multi-agent framework across its proprietary Canvas platform, affecting 130,000 professionals handling 160,000 audit engagements.
- Tax platforms: Deploying cloud-based document intelligence tools automated data extraction, reducing manual workloads by up to 90%.

The new initiative relies on shared governance, shared commercial models, and mutual accountability between both corporations to ensure practical execution for external corporate clients. ■



The alliance signals enterprise AI's shift from experimentation to organization-wide deployment with measurable business impact.



Why Google Antigravity 2.0 signals towards a multi-agent microservices future

Google's Antigravity 2.0 marks a shift from monolithic AI models to orchestrated multi-agent systems, redefining how enterprises build and govern AI at scale

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

Google launched Antigravity 2.0, expanding the platform from a specialized code editor into a standalone desktop operating layer. With this launch, Google is formalizing the transition away from monolithic artificial intelligence architecture.

For the past few years, Chief Information Officers (CIOs) have poured capital into massive large language models (LLMs) under a simple assump-

tion: bigger parameter sizes and larger context windows would yield better enterprise intelligence. However, when a single monolithic model tries to ingest an entire enterprise codebase, a massive data center log, or a complex financial structure, performance drops. Long-context degradation sets in, execution slowdowns occur, and the system begins to hallucinate under heavy operational load.

Beyond the monolith LLM

Traditional enterprise AI strategies relied on forcing a single model to act as a general practitioner—handling data retrieval, logical planning, code execution, and user interaction simultaneously. Antigravity 2.0 discards this approach through an agentic harness co-optimized with the newly released Gemini 3.5 Flash model. The architecture focuses entirely on the delegation of work.

Rather than crowding a single context window with a massive array of instructions, a primary agent defines and deploys dynamic subagents to execute narrow, specialized subtasks in parallel. These subagents process their workloads independently in the background, generate verifiable deliverables called artifacts, and report their refined findings back to the main agent.


Google demonstrated the raw capacity of this multi-agent microservices approach on stage. To test the system, engineers tasked Antigravity 2.0 with building a new operating system from scratch. The platform automatically deployed 93 separate subagents to process tasks in parallel, generating 2.6 billion tokens. The team completed the core framework in 12 hours for less than US \$1,000 in token costs, proving that orchestrated micro-agents can outmaneuver a single giant model.

Unattended automation and enterprise governance

For CIOs, the implications of Antigravity 2.0 extend far beyond a cleaner workspace. The platform detaches the AI from the integrated development environment (IDE) and introduces cron-like scheduled tasks, moving AI from a passive assistant to an active background operator.

Enterprises can now schedule independent agent teams to execute recurring, multi-step operations without human intervention. This shifts the executive focus toward establishing clear operational boundaries and robust compliance. For example, a team can schedule agents to audit cloud deployment logs hourly or review system analytics every morning automatically.

The introduction of JSON hooks also provides IT administrators with a concrete mechanism to control agent behaviors. Security teams can write simple JSON configurations to restrict access, intercept tool calls, or require manual validation for high-risk actions, such as executing external JavaScript code within the integrated browser. Furthermore, enterprise teams can log in using Cloud OAuth to route these workloads through



Antigravity 2.0 signals AI's next chapter; orchestrated micro-agents replacing monolithic LLMs for faster, scalable enterprise automation.

dedicated regional cloud endpoints, ensuring data stays within local compliance boundaries.

The agentic future

The structural design of Antigravity 2.0 signals the end of the enterprise chase for the largest parameter models to solve complex operational challenges. The massive, all-knowing single LLM is proving too slow and too costly for production-scale corporate pipelines.

Instead, the future belongs to the orchestration of smaller, hyper-specialized autonomous agents operating within a unified mesh. By deploying targeted subagents that collaborate under a strict governance layer, IT leaders can scale their digital operations sustainably. These new tools can potentially transform the role of CIOs from a manager of human software engineers to an enterprise architect of coordinated digital workforces. ■



THE NEW CIO

How AI agents, autonomous workflows and intelligent decision-making are reshaping the role of the CIO.

By Punam Singh | punam.singh@timesgroup.com

There is a question quietly circulating in boardrooms and executive corridors that few leaders are asking openly: What happens to an organisation when AI no longer just supports work, but begins to perform it?

For decades, enterprises were built around a simple assumption: people executed work, technology enabled it, and leaders managed both. That assumption is now beginning to fracture. AI is no longer confined to answering questions or automating isolated tasks. It is writing code, analysing contracts, making recommendations, coordinating workflows, and increasingly acting as an autonomous participant in business operations.

The implications extend far beyond technology. They challenge how organisations are structured, how decisions are made, how work is governed, and ultimately what leadership itself looks like. No executive sits closer to that transformation than the CIO.

The structure is already moving

Enterprises were constructed on a foundational assumption that has quietly collapsed: that humans are the primary executors of technology work. For twenty years, that assumption held. It built careers, departments, delivery frameworks, and entire outsourcing industries. It created the pyramid — wide at the base with executors at every level, narrow at the top where the strategists sat.

That pyramid is now being hollowed from the base upward — and the pace is not gradual.

Gartner projects that by 2030, 0% of IT work will be done by humans without AI involvement. The more immediate signal: 66% of enterprises are already reducing entry-level hiring as they deploy AI, according to IDC's 2025 Employee Experience survey, while 91% report that at least some roles have been changed or partially automated. The jobs disappearing first are precisely the ones the IT pyramid was built on — manual testing, tier-one support, infrastructure monitoring, routine maintenance.



The financial logic is blunt. Enterprise AI spending reached US \$37 billion in 2025 — more than triple the 2024 figure of US \$11.5 billion. That capital is not being deployed to augment existing teams. It is replacing the work those teams were doing. And yet most organisations remain staffed and evaluated as if none of this has changed. Job descriptions still list technical certifications as primary requirements. Performance frameworks still measure outputs rather than outcomes.

The gap between what the data shows and how organisations are responding is where the real risk lives.

The pattern is visible across every metric that touches the IT workforce — and the table below makes clear that this is not a single trend but several overlapping ones, each reinforcing the others.

The roles that produce employable juniors are contracting roughly twice as fast as the organisation is building new pathways to develop them. That is not a hiring slowdown. It is a structural break in how IT has historically grown its own talent. The two curves — entry-level hiring falling as AI-related hiring climbs — are now moving in visibly opposite directions, and the gap between them is the workforce problem every CIO in this story is contending with in some form.

The ROI trap that no one is discussing

Here is where the urgency sharpens: organisations are cutting headcount without capturing the returns they expected, and they are also paying more for the talent they are keeping — while the underlying cost structure of running AI operations is changing in ways most IT budgets are not yet designed to track.

Gartner surveyed 350 global business executives in 2025 and found that workforce reduction rates were nearly equal among those reporting higher ROI from autonomous technologies and those reporting modest gains or negative

"If I hired someone new, I wouldn't give them access to every system with full admin rights. We need to apply that same thinking to AI."



Mike Anderson
Global CIO,
Netskope

"Technology choices are now business choices."



Rajeev Soota
Vice President
— IT, Usha
International

outcomes. Around 80% of organisations report workforce reductions — but those reductions are not translating into measurable financial return. As Gartner Distinguished VP Analyst Helen Poitevin put it: "Workforce reductions may create budget room, but they do not create return."

At the same time, AI roles now command compensation three to four times higher than the average IT worker, and technical skill life cycles have compressed from 8 to 12 years to as little as 2 to 5. Organisations that overhire for AI talent risk locking into fixed cost structures that outpace business value. Those that underinvest find themselves unable to govern what they have already deployed.

Chitti Babu, Group CIO, Aurobindo Pharma, points to a structural mismatch that is quietly undermining AI ROI calculations across the industry. "Most enterprise IT budgets were built for a different world — fixed licensing plus headcount, capital expenditure on hardware, relatively predictable annual cycles. AI operations do not fit that model at all. You are now dealing with variable compute costs driven by GPU-hours and model size, inference costs that scale with requests per second, data ingestion volumes that compound as you connect more systems, and ongoing governance and compliance tooling that carries its own cost line. The budget model has to change before the AI strategy can be evaluated honestly."

The numbers behind his observation are significant. Data volumes feeding enterprise AI systems are projected to grow 2.5 times by 2027, driven by IoT, digital assets, and customer interaction data. The shift to larger language and multimodal models requires three to four times the compute for training compared to earlier AI generations. Roughly 70% of AI workloads are expected to migrate to managed cloud services — reducing capital expenditure but converting it into consumption-based spending that can spike unpredictably. Senior AI talent commands 20 to 30% salary premiums with no sign of stabilisation. Taken together, Babu's assessment is that AI operations costs are likely to roughly double by 2027 for enterprises that scale meaningfully — even as the work being done grows more capable.

"The CIO who goes to the CFO with a headcount-reduction story and no consumption cost model is setting a trap for themselves," he said. "You save 15 people on one side of the ledger and then discover you have a runaway compute bill on the other side that no one approved because no one designed the governance for it. Token consumption, inference costs, model retraining

cycles — these need to be tracked in real time and tied directly to business outcomes. Otherwise the CFO comes back in eighteen months and asks what we got for all of this. And if you do not have the answer, the AI budget becomes the first thing on the cutting-room floor."

Ashok Ramchandra Jade, Global CIO at Kirlskar Brothers shared his experience in scaling AI cost models and it echoes Babu's warning and sharpens it into something closer to a planning rule. "AI spend behaves like a utility bill, not a licence fee. The mistake most CIOs make is calculating cost during the pilot phase and assuming that number holds at scale — it never does. You have to plan capacity separately for each use case: tokens per task for a coding assistant versus a customer-facing chatbot, images or pages processed for a document pipeline, and then layer in users, tasks per user, and year-over-year growth, learned from the pilot itself." His rule of thumb is blunt enough to budget against directly: "Expect roughly double the cost in year two, and close to four times the pilot-year cost by year three, once usage scales properly. What looks lucrative during a pilot can look very different once it is live across the organisation — and if there is no clear, agreed benefits case with the business function footing the bill, that is the project that gets cancelled first when the bill arrives."

Meanwhile, Gartner predicts more than 40% of agentic AI projects will be cancelled by end of 2027 — driven by escalating costs, unclear business value, and inadequate risk controls. This is not a warning about future mismanagement. These projects are being planned and funded right now, in organisations that have not yet redesigned their governance, their talent mix, or the cost model underneath it all.

The CIOs navigating this moment well have accepted one uncomfortable truth: organizations of 2027 will not look like the enterprises of 2024

with AI tools placed on top of it. It will be structurally different, oriented differently, and measured against an entirely different definition of value.

The assumption that no longer holds

The traditional IT department was built for execution. Developers wrote code. Testers ran structured cycles. Infrastructure teams managed uptime. Support staff resolved queries across tiers. Scale meant headcount. The model was not inefficient for its time — it was precisely designed for the world that existed.

That world no longer exists.

AI writes production-grade code competently enough to pass review. It runs test scenarios faster than a QA team can design them. It monitors infrastructure around the clock, flags anomalies before a human notices them, and generates incident reports before an engineer has opened their laptop. And the scale of adoption is accelerating. Gartner forecasts that 40% of enterprise applications will be integrated with task-specific AI agents by the end of 2026 — up from less than 5% in 2025. By 2029, 70% of enterprises will deploy agentic AI as part of IT infrastructure operations, up from less than 5% today.

But the story looks meaningfully different depending on where you sit. In a technology company, the

shift from execution to orchestration is largely an internal challenge. In a manufacturing-led business with a dealer network across thousands of locations, legacy equipment on shop floors, and data living in systems never designed to speak to each other — the challenge is simultaneously more complex and more consequential. And the consequences of getting it wrong are not theoretical. Gartner warns that by end of 2026, "death by AI" legal claims will exceed 2,000 due to insufficient AI risk guardrails — a number that will only grow as autonomous decision-making embeds itself deeper into core operations.

The CIO who studies the business first

Before any structural redesign can happen, something more fundamental has to shift: how the CIO understands their own role.

Mike Anderson, Global CIO at Netskope, starts not with technology but with orientation. "As a CIO, you must be a student of the business that you're in. The only way you really do that is to go spend time with your frontline teams that are facing off to your customers. You can't change what you don't first understand. Intellectual curiosity has always been important. It's extremely important now with AI and how that's taking the role."

Cost component	Legacy IT budget model	AI operations cost model
Licensing / compute	Fixed annual licensing	Variable — tokens, GPU-hours, agent runtime
Infrastructure	Capital expenditure, on-prem	Cloud-native consumption, pay-as-you-go
Headcount	Fixed FTE salary cost	Premium AI talent, 20–30% above market
Data	Predictable storage growth	2.5× volume growth by 2027 (IoT, digital, interaction data)
Compute for models	Stable across cycles	3–4× higher for large/multimodal models
Cost predictability	Annual, budgetable	Spikes year-over-year as usage scales

Rajeev Soota, Vice President – IT, Usha International, describes the shift in terms of where the CIO's gravitational pull has moved. "AI has quietly but fundamentally changed what it means to be a CIO — especially in Indian and regulated enterprises. The role today feels far less about managing technology and far more about shaping how the organisation thinks, decides, and manages risk. With AI influencing business decisions directly, the CIO ends up sitting much closer to the core business agenda, working with the CEO, CFO, and business heads on growth, efficiency, and resilience. Technology choices are now business choices."

That last sentence is the most concise summary of what has changed. A decision about which AI system governs a procurement workflow or a credit assessment process is simultaneously a governance decision, a regulatory decision, a risk decision, and a business strategy decision. The CIO who brings only technology expertise to that conversation is bringing half of what the situation requires.

For Anjaiah Surgi, CIO, Hindware — the conglomerate behind Hindware — this business-first orientation plays out daily in board conversations about plant efficiency, demand forecast accuracy, and supply chain intelligence. "Traditionally, IT used to be called EDP — electronic data processing. It moved through infrastructure, then application management, then tech enablement. Now, with AI and ML, it is helping to define and redefine strategy. The way we approach the customer, the consumer, the partner ecosystem is changing. That is where the role of the CIO is evolving heavily."

AI at the heart of the enterprise

Much of the conversation around AI is still dominated by examples from the technology sector. But most economies run in manufacturing plants, dealer networks, and industries where the product is physical

"Demand forecasting used to drop as low as 45 percent. With AI, we are touching 85 to 90 percent. Just imagine the delta."



Anjaiah Surgi
CIO, Hindware

"Get governance right before scale forces your hand."



AJ Sunder
CPO/CIO & Co-founder,
Responsive

and the supply chain is complex. Surgi's organisation operates in exactly that world — and the results are measurable.

"Defect detection used to be manual. We implemented computer vision-based detection in the line almost two years back, accuracy is more than 98 percent, much better than the human eye. And demand forecasting in our sanitaryware business used to be around 60 to 65 percent accurate, sometimes dropping to 45 percent. Now, with AI-enabled forecasting, we are touching 85 to 90 percent. Just imagine the delta. The inventory turnover implications, the production capacity implications, the ability to put the right product in front of the right customer, end to end, it changes everything."

That delta — from 55% to 90%

forecast accuracy — translates directly into reduced working capital, fewer wrong-product production runs, and better dealer service levels. It is the kind of number that changes the CIO's seat in the budget conversation from cost justification to value demonstration.

This is what genuine AI ROI looks like at the operational level. It is also why the organisations still running fragmented, disconnected AI initiatives are falling further behind than headline adoption rates suggest. McKinsey's 2025 State of AI survey found that 88% of organisations now use AI in at least one business function, yet only 39% report any EBIT impact at the enterprise level, and just 6% qualify as genuine AI high performers. Nearly two-thirds of organisations say they have not yet begun scaling AI across the enterprise. Experimentation is widespread. Operationalisation is rare.

The new structure: three bands, not five layers

The traditional layer-cake org chart is giving way to something leaner. The emerging structure clusters IT talent into three distinct bands, defined not by technology function but by the nature of the work.

The first is the orchestration layer — people who manage, direct, and govern AI systems rather than execute what those systems perform. They set parameters, evaluate outputs, and intervene when something goes wrong. This profile barely existed five years ago. Today it is one of the fastest-growing talent demands in enterprise IT, and it commands a salary premium of roughly between 56% and 67% over traditional software roles.

The second is the judgment and domain layer — technologists who carry deep business context. Their value is not in building or running systems, but in interpreting what AI produces, challenging it when wrong, and translating capability into decisions that hold under scrutiny from a CFO or regulator.

The third is a smaller, highly specialised architecture and engineering layer — people who design the data foundations, integration frameworks, governance infrastructure, and security canvases on which AI systems run. This group is not shrinking. The cost of getting these decisions wrong compounds over time.

AI Sunder, CIO and Co-founder at Responsive, puts the logic plainly: "Execution was never the point — outcomes were. AI taking over repeatable execution frees the organisation to focus on what machines still can't do well: understanding business context, making judgment calls under ambiguity, and designing systems people actually trust. Our teams are shifting from building and maintaining to curating, governing, and connecting. The most valuable people in IT today understand both the technology and the business problem deeply enough to know where AI should be trusted, where it needs guardrails, and where a human needs to stay in the loop."

The talent challenge embedded in this restructuring is harder than it appears. IDC data shows that over 90% of global enterprises will face critical AI skills shortages by 2026, with those gaps putting up to \$5.5 trillion of economic value at risk. But the shortage is not evenly distributed. Demand for AI governance and orchestration skills is outpacing supply sharply — AI-related job postings grew 3.5 times between 2023 and 2025, with average time to fill AI roles rising from 42 days to 68 days. The organisations that redesign their IT talent model earliest will simply have more options.

The governance question that cannot be deferred

Where does AI-driven decision-making end, and where must human accountability begin? This is not a philosophical question. It carries operational, legal, and reputational weight — and the clock is running.

Anderson describes a framework that goes beyond standard "human



in the loop" language. "Every decision you make is a risk-based decision. Taking your highest risk decisions and trusting AI to be the decision maker is the wrong answer. How do we continuously adapt the trust of agents to take on more work — the same way we would with an employee? If I hired someone new, I wouldn't give them access to every system with full admin rights. We need to apply that same thinking to AI."

Anderson also named a failure mode that receives too little attention: human fatigue within automated workflows. "Machines work very quickly and humans work at a different pace. In a fully human workflow, if I got five requests today and suddenly I'm getting five hundred, are people rubber-stamping things? If we just go fast and don't think strategically, we'll end up with applications highly optimised for human workflows but not built for the agentic world. That's the biggest risk we run right now."

The regulatory dimension is tightening simultaneously. Soota at Usha

International points to accountability expectations that have fundamentally shifted. "AI has raised the bar on accountability. Earlier, system failures were operational issues. Today, automated decisions in finance, supply chain, HR, or customer experience can have regulatory or reputational impact. In the Indian context, data and trust matter more than speed. Regulations like DPDPA mean AI adoption cannot be reckless. CIOs are increasingly becoming custodians of data governance, consent, explainability, and compliance — not by choice, but by necessity."

Vinod Krishnan, CIO at Birlasoft, approaches governance through outcomes rather than process steps — a practical discipline born of resource reality. "I am not resourced to shadow the agent through every step. What I am resourced to do is look at the outcome, look at how it fits with standard outcomes I know, identify deviations, and stop them." His negative list model is blunt in its priorities: "Those activities I don't want anybody to do without super-

vision — fire people, spend money, send reports to government authorities, send reports to board members. Any outcome that touches one of these, I will see with my eyes before it goes."

Rakesh Bhardwaj, Global CIO, Lupin, brings the cross-functional model that regulated environments require. "Governance must go beyond data privacy to address accountability, ethical decision-making, transparency, and auditability. We focus on four essentials: responsible AI frameworks; cross-functional governance with legal, compliance, business, and technology leaders jointly overseeing AI from design to deployment; continuous monitoring and auditability; and building capability internally first. Decisions involving strategy, ethics, patient impact, regulatory judgment, or irreversible risk must remain human calls."

Ashish Kothari, VP and Partner at Kyndryl Consult India, flags the gap that sits beneath most governance conversations. "While ambition has accelerated, enterprise readiness has not always kept pace. CIOs are navigating the gap between AI experimentation and scalable execution. Integrating AI into legacy environments while managing risks such as data exposure makes governance and trust central to the CIO's mandate — particularly in India, where scale, cost sensitivity, and regulatory considerations demand a balanced, engineering-led approach."

Nikhil Prabhakar, CIO at IndiaMART, offers one of the more practical answers to where the human actually belongs in the loop — and his answer is that there is no single answer. "Human involvement depends entirely on the use case. Some AI workflows are fully automated — our IM VANI lead-handling conversations are an example where automation can operate independently without requiring human intervention. Other use cases are recommendation-oriented: AI generates intelligence, identifies clusters, segments users, and recommends

"You have to think about AI like a heart transplant, push it in such a way that the body accepts it, till one day it says: now this is mine."



Vinod Krishnan
CIO, Birlasoft

"Decisions involving strategy, ethics, patient impact, regulatory judgment, or irreversible risk must remain human calls."



Rakesh Bhardwaj
Global CIO, Lupin

actions, but a human decides what to do with that. A third category is genuinely human-in-the-loop — for meeting scheduling and productivity workflows, AI performs the preliminary work before handing the process to a human operator." His point cuts against the instinct to apply one governance posture across the whole organisation. "Different combinations of automation, recommendation, and human oversight are appropriate depending on business

requirements, return on investment, and accuracy expectations. Treating every use case the same way — either too cautious or too automated — is its own kind of risk."

The reskilling trap — and the talent paradox beneath it

Ask any CIO about their talent response and the answer is almost always the same: reskilling programmes, learning academies, internal bootcamps. None of this is wrong. Most of it is answering the wrong question.

Only 35% of leaders report having a mature, organisation-wide AI upskilling programme, according to recent research. Most training is fragmented, optional, and disconnected from the actual job tasks being transformed. IDC finds that 40% of IT leaders struggle with inconsistent skills development across their organisations, leaving them unable to measure true AI readiness. Teaching a manual tester to use an AI-assisted testing tool extends their employment for another cycle. It does not change what they are fundamentally employed to do — and it does not resolve the structural mismatch accumulating underneath.

Chitti Babu, Group CIO, Aurobindo Pharma, names this as a talent paradox rather than a talent shortage — a distinction that matters more than it first appears. "We are not running out of senior talent. The definition of a senior role is changing. On one side, AI is absorbing the routine work that used to form the entry point into IT — basic support, log analysis, first-pass code review. On the other side, we have a genuine and growing shortage of people who can design, govern, and interpret AI systems. That gap is not closed by a reskilling bootcamp. It is closed by intentionally rebuilding career pathways from the ground up."

The data makes the paradox concrete. Entry-level IT roles are contracting — Tier-1 helpdesk positions have declined as intelligent

agents resolve the majority of routine tickets without human involvement — while demand for AI governance, MLOps, and orchestration skills has surged to the point where AI-related roles now take an average of 68 days to fill, up from 42 days in 2023. The IT pipeline that used to replenish senior talent by graduating juniors through structured exposure to real systems is thinning precisely at the moment that senior AI expertise is most scarce.

Babu's prescription is architectural, not programmatic. "The answer is not more courses. It is rebuilding the junior role itself — reframing it around higher-value tasks. Data labelling, model validation, AI-ops monitoring, prompt engineering at the edge of a business domain. These are not consolation prizes for the roles that AI has taken. They are the new entry points into a career that leads somewhere. And alongside that, we need mentorship structures where juniors work directly with senior leaders on model training, data science problems, and AI ethics — learning by doing, not by watching a module."

This matters in regulated industries especially. At Aurobindo, where an AI-influenced decision about drug formulation or batch release carries regulatory and patient-safety consequences, the quality of the human judgment layer sitting above the AI system is not a secondary concern — it is the primary one. The senior talent shortage in those environments is not an HR problem. It is an operational and compliance risk.

Ashok Jade, another senior CIO voice in this conversation, pushes back gently on the premise that the profession is heading toward a hard senior-talent scarcity at all. His argument rests on precedent rather than prediction. "History shows technology changes jobs more often than it eliminates career ladders. The calculator did not eliminate mathematicians. Industry 4.0 did not eliminate manufacturing engineers. CAD and CAM did not eliminate civil engineers.

What changes is what the early years of a career look like — and that is exactly what is happening now."

His diagnosis of the actual risk is sharper than a simple shortage

"Governance and trust are central to the CIO's mandate — particularly in India, where scale, cost sensitivity, and regulatory considerations demand a balanced, engineering-led approach."



Ashish Kothari
VP & Partner,
Kyndryl Consult
India

"The CIO who goes to the CFO with a headcount-reduction story and no consumption cost model is setting a trap for themselves."



Chitti Babu
Group CIO,
Aurobindo
Pharma

narrative. "The real problem is not a lack of jobs for junior people. It is how they get practical exposure to what 'wrong' looks like, when AI is the one writing the code. A junior engineer used to learn what software quality means by fixing bugs, troubleshooting incidents, investigating low-risk security alerts. If AI is doing that work, juniors may never develop an eye for what a flawed system actually looks like — and that instinct is exactly what makes someone promotable into a senior role ten years later."

Jade maps the shift as a change in career architecture rather than its disappearance. Where the traditional path ran junior engineer to engineer to senior engineer to architect, the path now forming runs through AI-assisted system analysis toward AI orchestration — roles built around understanding systems deeply, recognising when AI is wrong, and managing increasingly autonomous environments rather than performing every task manually. "Companies need to deliberately design for this," he said. "One senior expert leading a team of junior developers, giving them as much hands-on exposure to real failures and real trade-offs as possible. That is how you create the next generation of senior people — not by hoping AI leaves enough manual work behind for them to learn from." His caution is narrower than a full talent-shortage warning but no less real: organisations reducing entry-level opportunities faster than they redesign training pathways should expect a temporary but genuine shortage of mid-level talent in the years ahead.

Sunder at Responsive frames the institutional dimension: "Governance and talent are two sides of the same coin. Agentic AI raises the stakes because the system is no longer just recommending — it's acting. That requires governance frameworks designed not just for outputs, but for autonomous behaviour. We invest in reskilling teams to think critically about AI-driven decisions and design workflows where human



“The real problem is not a lack of jobs for junior people. It is how they get practical exposure to what 'wrong' looks like, when AI is the one writing the code.”



Ashok Jade
Global CIO,
Kirloskar
Brothers

judgment and AI work together. We don't want AI fluency to sit in a small expert team. It needs to be distributed across the organisation.”

Krishnan at Birlasoft brought the most grounded perspective on why scaling AI fails — and it has nothing to do with data or infrastructure. “It is just the basic resistance that comes from introducing technology into the process of a company. The computer is deciding, but I am responsible — that is the immediate reaction. People fought Excel. People fought ERP. Today you try to take Excel away from an accountant, they will kill you first. You have to think about it like a heart transplant. You have to push it in such a way that the body accepts it, nurse that patient through the recovery cycle, give them medicines to lower the resistance. Till one day the body crosses that point and says: now this is mine.”

The human signal in the data supports his caution. Mercer's Global Talent Trends 2026 survey of nearly 12,000 executives and employees found that employee 'thriving' has collapsed from 66% in 2024 to 44%

in 2026 — below even pandemic-era levels — while worker concern about AI-driven job loss has risen from 28% to 40% over the same period. Organisations are deploying autonomous systems into a workforce that is increasingly anxious about them. The technical rollout is outrunning the human management of it.

The vendor evaluation rewrite

Krishnan articulated a vendor framework built on hard experience. “The first thing you want from a partner is basic competence — table stakes. Then comes diversity of experience. Third is your viability as a partner — a lot of the time the biggest risk with exciting new products is they don't stay in business long enough. And then your ability to go beyond the brief: to decompose what I ask for into what I want, and from what I want into what I need. Domain knowledge is central. The partner who speaks your business language — who talks about your assembly line, your defects per million, not IT terms — reduces your fear that they

“The 'I' in CIO is increasingly shifting from Information toward Intelligence and Innovation.”



Nikhil Prabhakar
CIO, IndiaMART

have misunderstood what matters.”

Surgi added the dimension specific to legacy-core environments. “Data integration capability is non-negotiable. My core is SAP ERP. Anything in connected systems must be properly integrated. If it does not integrate cleanly, data quality suffers, and if data quality suffers, AI decisions will suffer.”

Anderson flagged a risk the current AI application explosion has created — and it is a structural one.

With AI agent software spending forecast to reach \$206.5 billion in 2026 and \$376.3 billion in 2027 (up from \$86.4 billion in 2025), the vendor landscape is expanding faster than enterprise procurement teams can evaluate it. "The vendor ecosystem is going to get a lot more crowded because three people with an idea can build an application over a weekend. They're going to talk to someone in any function and show them a cool tool that didn't exist on Friday — but to them it looks like it has been around for five years because the website looks great. The reality is that was a great idea built with AI over the weekend. Did they have the right security posture behind it? Yet these companies thoroughly before you bring them into your environments."

The CIO of 2027: what the role actually becomes

Sunder offered the clearest articulation of where the role lands: "Two to three years from now, the CIO's most important function will be ensuring that an organisation's collective knowledge — its people, processes, data, and decisions — is structured in a way that AI can act on reliably and that humans can trust. That sits squarely between technology and business, which is exactly where the CIO role needs to live. Get governance right before scale forces your hand. Deploy AI in core workflows now and build institutional knowledge. Invest in people who can think across domains — that intersection is where the next generation of competitive advantage will be built."

Anderson extended this with a reframe of what the "I" in CIO might come to mean. "Chief integration officer, chief innovation officer, chief investment officer — or even chief risk and innovation officer. By 2030, organisations are going to be asking how we measure the effectiveness of token consumption because there is a lot of investment in AI and not a lot of thought around what it is costing. Someone has to be the custodian

of that. I also spend as much time with our Chief People Officer as I do with our CISO today — talking about human risk, who are the intellectually curious people tinkering with AI, and how do we amplify their skills. There is a lot of reverse mentoring that can go on if you are willing to accept it."

Nikhil Prabhakar, CIO at IndiaMART, arrives at a similar reframe from a different direction — tracing how the letter has already moved once before it moves again. "Historically, the 'I' represented information, because technology teams were primarily responsible for providing data, reports, forms, and ERP systems. That role has already evolved significantly. Today, technology organisations are increasingly responsible for delivering business intelligence rather than simply information, and over time CIOs will become even more accountable for business KPIs themselves. The role is moving from information toward intelligence and innovation." His view of what that means in practice is concrete. "Future CIOs will focus less on managing systems and more on customer experience, business strategy, productivity, technology-enabled growth, and measurable business outcomes. The objective is no longer simply maintaining infrastructure — it is enabling and managing the performance of the entire business ecosystem. That is why I believe the 'I' in CIO is increasingly shifting from information toward intelligence and innovation."

The decision every CIO is making

Strip away the governance frameworks, the org chart redesigns, and the vendor criteria, and the fundamental decision every CIO is navigating comes down to this: how much of the organisation's decision-making are you willing to delegate to a system you cannot fully audit, in exchange for the speed and scale it provides?

There is no objectively correct answer. The right balance depends

on the decisions involved, the maturity of the AI systems, the regulatory environment, and the trust the workforce has developed in the technology over time.

What is clear is that organisations deferring this decision are accumulating risk faster than they realise. Every automated workflow without a clear accountability framework, every AI output that goes unreviewed because the review process was never designed, every reskilling programme that teaches tool usage without redesigning the underlying role — each represents a debt that will eventually be paid. Gartner predicts that 50% of organisations will be forced to implement "AI-free" skills assessments by end of 2026 because generative AI use has visibly degraded critical-thinking skills. That is a governance failure showing up in a talent metric.

The CIOs in this story are getting ahead of that debt. They are redesigning organisations before structural gaps become crises, establishing governance before regulators demand it, redefining talent around judgment rather than execution, and measuring AI not by deployment, but by business outcomes. As Surgi puts it, without clear success metrics, nothing can be sustained, regardless of how well it begins..

AI is no longer just another technology initiative. It is becoming part of the enterprise itself. As humans and intelligent agents begin working together across functions, the competitive advantage will shift from deploying AI to governing it well—embedding it into business processes, aligning it with organisational goals, and earning the trust of employees, customers, and regulators alike. The CIO's role is expanding accordingly: from managing technology to shaping how the enterprise operates, decides, and grows. In the years ahead, the most successful CIOs won't simply lead digital transformation. They will help define the future of the enterprise. ■

COMPUTEX 2026: From Personal Devices to AI Infrastructure



AI infrastructure dominated COMPUTEX 2026, reshaping enterprise computing and consumer technology

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

For years, COMPUTEX Taipei was an exhibition where tech enthusiasts came to see flashy gaming rigs, watch overclocking contests, and explore the latest laptops and PC hardware. It was, above all, a show about personal technology.

That spirit is still there. But walking the halls of the 45th COMPUTEX, held from June 2–5, 2026, under the theme “AI Together,” it became clear that the industry’s focus has moved well beyond

PCs and consumer gadgets. The change has been visible for a few years now, but this year’s show felt different. AI wasn’t simply one of the big stories at COMPUTEX, it was the story, shaping conversations across nearly every booth and product announcement.

Now, companies attracting the most attention aren’t showing faster gaming laptops. Instead, they unveil server racks, AI accelerators, and liquid



This year, 111,312 buyers and visitors from 152 countries and regions, including Japan, the United States, South Korea, China, Singapore, India, and Vietnam, came to Taipei to see where the tech industry is headed

cooling systems for the next wave of artificial intelligence. In short, COMPUTEX has matured!

This year, 111,312 buyers and visitors from 152 countries and regions, including Japan, the United States, South Korea, China, Singapore, India, and Vietnam, came to Taipei to see where the tech industry is headed.

Key announcements

NVIDIA arguably drew the most attention at the event. The company unveiled RTX Spark, an ARM-based laptop processor developed with MediaTek, marking an entry into a market it has largely stayed away from.

For a company best known for its GPUs, the move reflects a broader



ambition: to have a presence wherever AI computing happens, including the laptops people use every day. It is another sign of how NVIDIA is expanding beyond its traditional roots. The Santa Clara, California-based company has become one of the biggest beneficiaries, and architects, of the AI era, providing the chips and computing platforms that power much of today's generative AI ecosystem.

It also introduced Vera Rubin, a new AI architecture designed for

autonomous AI systems moving from research labs into real businesses. Instead of just answering questions, this AI can act on its own and needs robust infrastructure to support it.

Network-Attached Storage (NAS) devices and software solutions provider Synology took a different but equally interesting approach. Its updated DiskStation Manager platform lets you use AI entirely on your own hardware, without needing the cloud. This means smart search,



automated workflows, translation, and admin help, all while keeping sensitive company data on-site. The company also introduced ActiveProtect Manager 2.0, which uses AI to find and isolate suspicious files before they cause problems, and its PAS7700 enterprise storage system, which holds 216 NVMe drives and delivers 30 GB/s of sequential throughput.

AMD and Qualcomm each shared their own ideas about the future of AI. AMD focused on its Ryzen AI processors, saying your PC should handle more tasks locally instead of sending everything to a data center. Qualcomm made a similar point, especially about energy savings, which matters for devices that need to run AI all day on a single charge. Together, they highlighted a future in which AI workloads are dynamically distributed across PCs, edge environments, and the cloud, balancing performance, cost, and energy efficiency.

Transformation of the global exhibition

Three years ago, at COMPUTEX 2024, everyone talked about AI PCs—laptops with special AI chips and even a Copilot key on the key-

board. Back then, it felt more like a novelty, something to show off rather than something that truly changed how you worked.

By 2025, AI in laptops, gaming handhelds, and accessories made a real difference, with better battery life, smarter performance, and features people used.

In 2026, the conversation shifted again, this time in a bigger way. Now, it's less about what AI can do for your laptop and more about what it takes to build AI systems that run whole companies. The hardware on display shows that ambition—not just personal devices, but the large-scale infrastructure behind them.

The exhibition floor tells the story

At COMPUTEX 2026, one could see server cabinets as tall as a room, complex liquid-cooling setups, and networking gear built to move data at speeds that seemed impossible a few years ago. Vendors weren't just showing single products, they presented full AI infrastructure stacks, from chips and cables to the software that connects everything.

NVIDIA described this as building "AI factories," meaning that future data centers won't just store and

process information—they'll create intelligence at scale. Intel brought its own story about edge AI and robotics. Many other vendors added innovations in cooling, storage, and connectivity, all of which are quietly essential for keeping these systems running.

It's very different from the days of overclocked gaming PCs. But in its own way, it's just as exciting if you know what to look for.

Gaming and consumer tech aren't going anywhere

Still, COMPUTEX hasn't forgotten its roots, and the crowds around gaming and consumer areas showed personal technology remains very important.

ASUS celebrated 20 years of its Republic of Gamers brand at the show, a major milestone. They introduced AI-powered gaming laptops, handheld devices, creator systems, and many peripherals, showing how much the ROG line has grown. AMD highlighted its gaming strengths and AI efforts. MSI, Acer, Gigabyte, and others kept the COMPUTEX tradition with many new laptops, monitors, graphics cards, and handheld gaming devices.

These consumer products are

One important trend at the show was the move toward edge AI. This means running AI locally on hardware you own and control, rather than relying solely on remote cloud servers.

no longer just standalone gadgets. Increasingly, they're the user-facing part of a bigger system, connected, sometimes visibly, sometimes not, to edge infrastructure, enterprise platforms, and cloud AI. Your gaming laptop is still your gaming laptop, but now it's part of something larger.

AI closer to home

One important trend at the show was the move toward edge AI. This means running AI locally on hardware you own and control, rather than relying solely on remote cloud servers.

This matters for several reasons. Privacy is one—organizations want to run AI on sensitive data without it leaving their building. Speed is another: local processing is faster than sending data to the cloud and back. Resilience is also important since systems that work without a constant internet connection are more reliable.

Global cybersecurity company Trend Micro showcased AI-powered security tools designed to protect the next generation of AI-driven businesses from increasingly sophisticated threats. Overall, AI is spreading from central data centers to the edges of networks and devices at their ends, a more distributed and resilient intelligence.

Robots are having their moment

If anything surprised people at the show, it was robotics.

These weren't just robots like in past tech shows. Exhibitors showed



real systems: autonomous warehouse robots, industrial manufacturing platforms, AI assistants for physical spaces, and smart logistics technologies. The key point was that this AI doesn't just think it acts in the real world with real results. In healthcare, manufacturing, logistics, and smart cities, these technologies had wide-ranging uses and, in many cases, were already being tested. In Taipei, it felt like physical AI was moving from demos to real-world use faster than most expected.

The startups worth watching

If you wanted a different kind of excitement, the InnoVEX startup pavilion was the spot to visit.

While big company booths had polished presentations and rehearsed demos, InnoVEX felt livelier and more chaotic. It was a place where you could talk directly to a founder and get a real answer about what their technology does and why it matters. The startups covered robotics, semiconductors, green tech, AI infrastructure, healthcare, and smart mobility.

This year's event had over 500 startups, an 11% increase from last year, from 23 countries. Taiwan, Japan, and South Korea had the larg-

est groups. For many visitors, this was the most valuable part of the show—a chance to see technologies that could shape COMPUTEX in the next five years.

So, what is COMPUTEX now?

COMPUTEX began as a show for computer parts. It grew into the world's top event for gaming hardware. Now, it's harder to sum up in one sentence: a global stage for AI infrastructure, enterprise computing, robotics, and intelligent systems, with consumer tech still beating beneath it all.

The RGB lighting is still there, and gaming rigs still attract crowds. But this year, the most important conversations happened in front of server racks, cooling systems, and robotics platforms that most people wouldn't have recognized five years ago.

If 2024 was about bringing AI to COMPUTEX, and 2025 about making it useful for everyday people, then 2026 feels like the year the industry focused on building essential, behind-the-scenes AI foundations that everything else relies on, even if they aren't glamorous and cost a lot.

—The author attended COMPUTEX 2026 at the invitation of Synology. ■



How Synology is reinventing itself for the AI era

Andrew Huang on AI sovereignty, data ownership, and why trusted infrastructure will define the next phase of enterprise AI

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

As enterprises race to deploy generative AI, attention is shifting from models and applications to a less glamorous but increasingly strategic layer of technology: data infrastructure.

Organizations are discovering that AI success depends not only on algorithms but also on how effectively they can store, govern, protect, and access growing volumes of data. Questions around where data resides, who controls it, and how securely it can be used are rapidly becoming boardroom discussions.

For Synology, a company long associated with network-attached storage (NAS), this shift

presents an opportunity to evolve beyond its traditional identity.

"Synology began as an SMB-focused storage provider, and over the years we have built a strong presence in that segment, including in India," says Andrew Huang, Regional Sales Manager at Synology, speaking on the sidelines of COMPUTEX 2026 in Taipei. "However, our focus today extends well beyond storage. Over the past three years, we have been making a major push into the enterprise market," he adds..

That shift reflects Synology's broader ambition to evolve from a NAS vendor into an enterprise

data management and infrastructure provider. The Taiwanese company, known for its network-attached storage (NAS) systems and on-premises data management solutions, is targeting aggressive growth in India and expects to double its business in the country over the next three years, driven largely by rising demand for enterprise-grade storage, data protection, and AI-ready infrastructure.

Synology's NAS boxes allow enterprises, small businesses and individual consumers to store data locally offering an alternative to cloud-based storage services.

Today, Synology's ambitions span data protection, surveillance, productivity, and AI-enabled infrastructure. Storage remains the foundation, but the company increasingly positions itself as a provider of trusted data infrastructure for the AI era.

From storage to AI infrastructure

The rapid adoption of AI is driving unprecedented demand for storage performance, scalability, and resilience. Every AI initiative depends on large volumes of data that must be collected, managed, protected, and accessed efficiently.

"Every AI initiative ultimately depends on a robust infrastructure layer, and storage sits at the center of that foundation," Huang explains.

To address these requirements, Synology recently introduced the PAS7700, its flagship dual-controller NVMe storage platform designed for AI workloads, analytics applications, and other data-intensive enterprise environments where performance and availability are critical.

But hardware is only part of the company's strategy.

Synology is increasingly investing in private AI environments that allow organizations to deploy AI capabilities within their own infrastructure rather than relying entirely on public cloud platforms.

"Unlike public AI services, which rely on cloud-hosted models and external data processing, our focus



Andrew Huang
Regional Sales Manager, Synology

is on helping organizations build AI environments within their own infrastructure," Huang says. "This enables enterprises to retain full control over their data while leveraging AI technologies securely and efficiently."

The company is also developing GPU-enabled capabilities that will allow enterprises to run AI models and training workloads on-premises using their own data assets.

The rise of AI sovereignty

As organizations experiment with generative AI, concerns around privacy, governance, compliance, and data ownership are becoming increasingly prominent.

Many enterprises remain cautious about sending sensitive corporate information to external AI platforms, particularly in highly regulated sectors such as banking, financial services, healthcare, government, and education.

For Synology, this challenge has evolved into a core strategic focus.

"Our approach revolves around what we call true AI sovereignty," says Huang. "Organizations should be able to centralize their digital assets in a secure environment and use AI capabilities without having to move their data outside their own infrastructure."

The concept aligns with a broader global trend. Governments and regulators are introducing stricter

requirements around data residency and privacy, while enterprises are demanding greater visibility into how AI systems access and process information.

For organizations operating under strict governance requirements, maintaining ownership and control of data is becoming non-negotiable.

"We believe the future of enterprise AI will be built around trust, ownership, and control," Huang says.

The AI-driven storage explosion

While public discussions around AI often focus on chatbots and productivity tools, Huang believes the larger impact will be felt at the infrastructure layer.

According to him, AI fundamentally changes enterprise data economics. Every AI workflow generates additional datasets, outputs, models, derived assets, and archived records, creating exponential growth in storage requirements.

"Every AI workflow involves collecting, processing, analyzing, and often creating additional data," he says. "In many cases, AI multiplies the volume of information organizations need to manage."

The next wave of AI adoption is also expected to move beyond text.

"Today, much of the conversation revolves around text-based AI, but the future will increasingly involve video, images, sensor data, and other unstructured content. These workloads require significantly larger storage capacities and much faster performance."

As a result, demand for high-performance storage infrastructure is expected to grow significantly. However, simply adding capacity will not solve the problem.

Organizations must also address increasingly complex questions around data governance, security, backup, recovery, and lifecycle management.

"As data volumes continue to grow, enterprises will need better ways to centralize, govern, and pro-



“Every AI initiative ultimately depends on a robust infrastructure layer, and storage sits at the center of that foundation.”

protect their information assets,” Huang says.

Growing focus on data protection

In India, Huang believes awareness around data protection remains uneven.

While large enterprises have made significant investments in cybersecurity, backup, disaster recovery, and compliance, many smaller organizations are still developing mature data protection strategies.

“One of the most significant challenges is awareness around data protection,” he says.

The challenge is becoming more important as India’s digital economy continues to expand and organizations face increasing scrutiny around data privacy and governance. The Digital Personal Data Protection (DPDP) framework is expected to further strengthen conversations around data management, security, and accountability.

As AI adoption accelerates, organizations of all sizes will need to view data protection as a business priority rather than simply an IT requirement.

“We work to help organizations

understand that data is one of their most valuable business assets,” Huang says. “Protecting it should be a strategic priority, not just an IT function.”

Building through partnerships

A key pillar of Synology’s growth strategy in India is its partner ecosystem.

“System integrators are a critical part of our growth strategy,” Huang says.

As enterprises navigate increasingly complex infrastructure environments, local partners often serve as trusted advisors, helping organizations design and implement modern storage, protection, and governance frameworks.

Synology is actively expanding its network of system integrators and channel partners across the country to strengthen its reach and support enterprise customers more effectively.

The company is placing particular emphasis on sectors such as manufacturing, government, BFSI, and education, where digital transformation initiatives are creating growing demand for secure and resilient data infrastructure.

The defining question for enterprise AI

Looking ahead, Huang believes the industry will increasingly be shaped by one simple question: How much do organizations trust their own data?

As enterprises move from AI experimentation to large-scale deployment, issues such as data quality, ownership, governance, and accountability will become as important as model performance.

“Organizations are realizing that AI is only as effective as the data behind it,” he says.

For Synology, that realization underpins its long-term investment strategy.

“Our biggest priority is AI sovereignty,” Huang says. “The goal is to help enterprises build AI capabilities without compromising control over their digital assets.”

If Huang’s thesis proves correct, the next competitive advantage in enterprise AI will not come solely from access to better models. It will come from an organization’s ability to build trusted, governed, and sovereign data environments that allow those models to operate securely and effectively.

In that future, infrastructure may become as important as intelligence itself. ■

The CEO agenda for 2026: AI, agility and workforce transformation



Employees are not just looking for direction – they are looking for context. What is changing, why it matters, and how it impacts them

By **Sunil Jose** | editor.tech@timesgroup.com

What is remarkable about the challenges that today's CEO faces is that it is no longer an issue of having one priority versus another, but rather how a number of changes are converging at once. Especially, the implementation of artificial intelligence, the increasing demands on the workforce, and the pressure to turn strategy into measurable outcomes.

For CEOs, this means the focus is no longer

on driving standalone initiatives. The real challenge is aligning technology, workforce strategy, and business outcomes in a way that is both practical and scalable so that AI delivers impact, not just experiments.

From pilots to integrating AI into decision-making

Across organisations, AI is already being deployed.

However, the future lies in making it purposeful. In most instances, the AI application is still isolated in pilots within functions or use cases, which do not align entirely with organisational goals. What we observe now is a transition towards integrating AI into decision-making processes.

Data plays a crucial role here. Without a connected perspective on people, financials, and operations, it would be hard to go from isolated insights to actions. Inefficiencies create delays when fast decisions are needed.

As the use of AI becomes more ingrained in the organisational fabric, leaders are becoming increasingly concerned with the importance of trust. They focus on the process of decision-making, the transparency of AI systems, and where human intervention is necessary, as well as the safeguards in place when AI is acting on critical business data. Organisations that succeed in this endeavour are not the ones utilising the most AI applications but rather those that leverage AI in a manner that serves clear business objectives.

Rethinking agility as a core capability

Agility is often spoken about, but in practice it is still evolving. Over the last few years, many organisations have responded well to disruption. The next step is building the ability to do this consistently, not just in moments of crisis.

This has triggered the shift from periodic to ongoing decision-making driven by data. Instead of following quarters or annual plans, executives use real-time indicators for guidance.

Having a connected view across functions plays an important role here. It helps leaders understand trade-offs in real time, whether that means reallocating budgets, shifting talent, or revisiting investment priorities.

Workforce transformation is already underway

Workforce transformation is often



Sunil Jose
President – India, Workday

discussed as a future trend. In reality, it is already happening. Roles are changing, skill requirements are evolving, and employees are looking for more clarity on how they grow within organisations.

In India, this is particularly visible. Organisations are balancing rapid digital adoption with the need to create more personalised and meaningful employee experiences. One of the shifts we are seeing is towards skills-based thinking. Instead of managing roles as fixed structures, organisations are beginning to look at capabilities more dynamically – what skills they have, what they need, and how they can build them over time. This is especially important as AI changes work itself. Organisations need clear insight into their skills today and a plan to build the skills they will need as AI becomes part of everyday processes.

Another important change is the use of real-time employee feedback. The ability to listen and respond in real time is becoming a key driver of both engagement and retention. These are not isolated changes. They reflect a broader move towards making workforce decisions more responsive and better aligned with business needs.

Bringing technology and talent together

What is becoming clearer is that

technology and workforce strategy cannot be addressed separately. AI initiatives depend on having the right skills in place. At the same time, workforce transformation relies on having the right systems to provide visibility and insight.

This puts a greater emphasis on how leadership teams work together. CIOs, CHROs, and business leaders need to operate more closely, with shared accountability for outcomes. From a CEO perspective, this is about ensuring that these efforts are aligned – not just in intent, but in execution. In our conversations with leaders across India, we see the most progress where this accountability is clearly shared, and AI is treated as a lever for business results, not just technology rollouts.

The importance of clarity

Clarity regarding AI implementation, new responsibilities, and organisational objectives is critical for establishing trust and alignment. On the other hand, it is equally crucial to remain realistic about areas of continuity. No matter how much technological advancement occurs, the importance of humans should never be underestimated.

The goal here is to leverage technology to enhance decision-making, user experience, and employee engagement.

Looking ahead

The 2026 CEO Agenda does not involve picking between AI, agility, and workforce transformation. Instead, it involves finding ways of combining the three in such a manner that they fit the organisation in question. This is not a transition but a process involving constant adaptation and learning on decision-making methods.

The organisations that will move forward with confidence are those that keep this approach practical – linking AI investments to clear outcomes, staying close to their people, and building the ability to respond to change as it happens. ■

Any AI strategy must be tightly connected to a data strategy

Robert Pizzari, Splunk's Asia head, explains how Cisco Data Fabric helps Indian enterprises prepare for AI while balancing compliance, cost, and resilience.

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

As enterprises accelerate AI adoption, the conversation around data infrastructure is rapidly shifting from experimentation to operational readiness. Organisations across India are now grappling with a mix of rising data volumes, increasing regulatory pressure, legacy infrastructure, and the growing need to operationalise AI securely and cost-effectively. At the same time, CIOs and CISOs are being pushed to rethink how machine data is collected, governed, and utilised across hybrid IT environments that span on-premises systems, multi-cloud architectures, and emerging AI workloads.

In this interview, Robert Pizzari, Group-Vice President, Asia- Splunk discusses how Splunk, backed by Cisco, is positioning itself beyond cybersecurity and into broader operational intelligence. He talks about the company's evolving data strategy through the Cisco Data Fabric, the growing importance of federated search, and how enterprises can make AI-ready data environments more cost-efficient without abandoning existing infrastructure investments. He also shares insights into India's regulatory landscape, the mounting pressures faced by CISOs, the role of machine data in the AI era, and why organisations need a tightly aligned AI and data strategy to remain resilient in an increasingly automated world.

CIO&Leader: Splunk is often seen as a security tool for enterprises, but you are also positioning yourself as an operational engine. How does Splunk bridge the gap between raw IT telemetry and board-ready business intelligence?

ROBERT PIZZARI: At the end of the day, our heritage over the last 20-plus years originated as a monitoring tool focused on finding a needle in a haystack across applications, infrastructure, databases, and other systems. We really have not moved too far away from that foundation.

We fell into security as opposed to specifically going after security because our real strength has always been as a data platform. One of the recent announcements following the Cisco acquisition is the concept of the Cisco Data Fabric. It is designed to give organisations the ability to move, access, and govern data regardless of where it lives across networks, clouds, applications, and now AI systems.

There is also another layer tied to the machine data lake. This is not limited to IT, applications, databases, or security specifically. The machine data lake capability helps organisations land all their machine data irrespective of where it resides and helps them make that data AI-ready from day one without forcing them to move or repackage data constantly across repositories.



“By decoupling the need to centralise all data, we believe we can help organisations maintain better control over cost, governance, and compliance without forcing them to move everything.”

Another important transition we have taken is around federated search. Historically, if you asked whether all data needed to sit inside Splunk, the answer would have been yes. Most solutions in the market today, including SIEM platforms, still require centralised data repositories.

By decoupling the need to centralise all data, we believe we can help organisations maintain better control over cost, governance, and compliance without forcing them to move everything into Splunk. This federated search journey started around two years ago with support for Amazon Web Services S3 buckets, and recently we also announced support for Snowflake. Over time, more data lakes will be supported as well.

CIO&Leader: AI-ready data often comes with storage and processing challenges. AI adoption in India is expensive for enterprises. How does Splunk help organisations maintain cost efficiency?

ROBERT PIZZARI: When I think about the mindset of a CIO or even a

CISO, it is really about doing things efficiently while balancing cost and operational performance.

As we move into the AI era, organisations are increasingly focused on securing data, monitoring AI stacks, and using AI to improve outcomes such as faster detections, quicker investigations, and better uptime. We estimate that around 1.3 billion AI agents will become active by 2028, while global data volumes could reach around 394 zettabytes. More than half of that volume is expected to be machine data.

That obviously creates enormous challenges around storage and processing costs.

With the Cisco Data Fabric, organisations can access and train data whether it sits inside Splunk or elsewhere. We also enable search to be applied dynamically depending on the use case. One important thing we have learned over the last two decades is that not all data is created equally.

Some data use cases are near real time and require optimisation, while others are more focused on compliance and regulatory requirements

and are accessed less frequently. By democratising the way data is accessed, we believe organisations can achieve much more cost-effective outcomes.

CIO&Leader: How should Indian CIOs rethink their approach to data collection? At what point does data become a liability instead of an asset?

ROBERT PIZZARI: Without doubt, any AI strategy must be tightly connected to a data strategy.

Right now, data is fragmented across multiple tools, repositories, lakes, rivers, and environments. If organisations do not have a strong data strategy, AI is not going to magically solve the problem. In fact, poorly orchestrated data can actually create more noise.

As we move into an agentic AI era, autonomous systems will increasingly depend on real-time, contextualised, high-quality data. The old principle of “garbage in, garbage out” still applies. AI systems are not automatically going to fix poor data hygiene or eliminate risks such as hallucinations and bias.

That is why machine data strategy is becoming so important. For example, if a CIO already invested heavily in platforms like Snowflake or Amazon over the last few years, we are not asking them to move all of that data into Splunk. They can continue leveraging those investments while using Splunk for operational and machine data use cases where time is critical.

As attackers increasingly adopt AI, organisations also need to rethink how they defend themselves. Traditional defense models are no longer enough. Machine data growth across endpoints and repositories means organisations will need normalised and streamlined data environments so they can operate efficiently with AI systems.

Data tiering is another critical area. Organisations need to decide what data requires real-time access and what can be archived or

accessed periodically. Our goal is to help organisations make those decisions while preserving the investments they have already made.

CIO&Leader: Indian organisations still rely heavily on legacy systems while also adopting modern cloud environments. How does Splunk support this hybrid reality without forcing organisations into complete cloud migration?

ROBERT PIZZARI: This is not unique to India, but it is definitely a major reality there. Most organisations that have existed for a long time still rely on legacy applications and architectures. Some of these systems may remain operational for decades because replacing or rebuilding them is extremely difficult and expensive.

With Splunk and Cisco's broader Data Fabric strategy, we support everything from on-premises legacy systems all the way to modern cloud-native environments. We are not asking organisations to abandon existing systems to benefit from innovation.

Take a typical bank as an example. Many banks operate thousands of applications, with core banking systems that may have been running for 10, 20, or even 30 years. Those systems are not going away anytime soon because replacing them would involve enormous cost and operational complexity.

For us, it does not matter whether the environment is legacy or future-state. Our monitoring capabilities are designed to support everything from on-premises systems to hybrid and cloud-only infrastructures.

Another important point specific to India is the regulatory landscape. CERT-In requires incident reporting within six hours, while the Digital Personal Data Protection Act introduces additional disclosure obligations. These timelines make broad monitoring strategies absolutely critical.

Even sectors like healthcare still operate operational technology sys-

“Any AI strategy must be tightly connected to a data strategy.”

tems running older operating systems that cannot easily be patched. Splunk is one of the few companies capable of supporting both older monolithic technologies and modern architectures such as Kubernetes and microservices.

CIO&Leader: What are your top recommendations for organisations trying to align AI strategy with data strategy?

ROBERT PIZZARI: The first recommendation is visibility. Organisations cannot manage what they cannot see. Most enterprises already know they have monitoring blind spots and operational silos that need to be addressed.

The second recommendation is simplification. Organisations should review how many monitoring tools they have accumulated over the years and evaluate whether some of those tools can be consolidated to improve efficiency, reduce costs, and simplify operations.

The third recommendation is building a strong foundational data strategy powered by the Cisco Data Fabric. We are not suggesting customers replace all of their current investments. One of Splunk's key strengths is that we continue to operate as an open ecosystem platform even after the Cisco acquisition.

For organisations already using Cisco technologies across networking, data centers, connected workplaces, and AI systems, we are improving native integrations to simplify data ingestion and reduce reliance on unsupported third-party integrations.

Another major part of our strategy is allowing organisations to bring their own AI models to Splunk data.

We are not building a closed environment where customers are restricted to our AI models. Through APIs and MCP servers, organisations can bring in foundational models, open-source models, or internally trained models depending on their governance requirements.

This is especially important for sectors like banking and government, where AI models may need to remain entirely within controlled on-premises environments. By reducing vendor lock-in and democratising access to data, organisations can scale more effectively while controlling costs.

CIO&Leader: Many organisations still view cybersecurity as a cost center. How does AI-powered automation help CISOs prove ROI and position security as a business enabler?

ROBERT PIZZARI: The key lies in being able to ingest data once, pay for it once, and use it across multiple teams and business functions.

For example, the same telemetry collected for cybersecurity purposes can also support business analytics and customer insights. Organisations can analyse customer behavior, purchasing patterns, usage trends, and other business metrics using the same underlying datasets.

This helps organisations avoid paying multiple times for the same data.

There is also a major difference between cybersecurity incidents and operational incidents. Cyber incidents in regulated industries come with compliance obligations and require forensic integrity and chain-of-custody processes. At the same time, operational teams may use the same datasets to improve customer experience or application performance.

Our overall approach is focused on simplifying monitoring, making data accessible regardless of where it resides, and enabling different teams to use that data in a governed and secure way. ■

Many companies are still building human-optimised applications

Mike Anderson, Global CIO at Netskope explains why future-ready enterprises must redesign governance, security, and workflows for an AI-native world built on trust, oversight, and adaptive intelligence.

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

Enterprise technology leadership is entering a phase where AI adoption is no longer limited to experimentation or productivity enhancement. Organisations are now confronting a deeper structural shift, one that affects software architecture, governance, workforce models, security frameworks, and enterprise decision-making itself. As agentic AI systems begin participating in workflows once handled entirely by humans, CIOs are being pushed beyond their traditional mandate of managing infrastructure and applications. They are increasingly expected to balance innovation with accountability, accelerate transformation while controlling risk, and align AI investments with measurable business outcomes.

In this interaction with CIO&Leaders, Mike Anderson discusses why the future enterprise will require a complete redesign of how humans, AI agents, and business systems interact. From adaptive trust and zero trust security to AI governance, human oversight, organisational fatigue, and the economics of AI adoption, he explains why the next generation of CIO leadership will depend less on technology administration and more on orchestrating resilient, intelligent, and trustworthy digital ecosystems.

CIO&Leader: The role of the CIO is changing rapidly with the rise of agentic AI. How do you see this transformation?

MIKE ANDERSON: A CIO must first become a student of the business they operate in. The only way to truly understand a business is by spending time with frontline teams, customers, and engineers building products for the market.

Technology leaders can no longer operate only as infrastructure managers. AI is forcing CIOs to become deeply involved in business transformation, operational strategy, workforce redesign, and investment prioritisation.

Intellectual curiosity has always mattered in technology leadership, but in the AI era it becomes critical. The pace of change is too rapid for static leadership models. CIOs must continuously learn, experiment, and understand how AI changes the mechanics of business itself.

CIO&Leader: AI is increasingly entering strategic decision-making processes. Where should organisations draw the line between AI recommendations and human judgment?

MIKE ANDERSON: Every enterprise decision is ultimately a risk-based decision. High-risk decisions should never be delegated entirely to AI systems.

Organisations need to build trust in AI progressively, much like they build trust in employees. A new employee is not immediately given unrestricted access to systems or authority. Responsibility expands gradually as confidence grows. The same principle should apply to AI agents.

AI can significantly improve areas like scenario planning because it enables organisations to evaluate multiple outcomes at machine speed. But human oversight remains essential, especially in areas involving financial commitments, governance, regulatory exposure, or operational risk.

The challenge is determining which decisions require direct human intervention and which can safely operate with supervised autonomy.

CIO&Leader: Netskope's research highlights disconnects within the C-suite around AI adoption. What are the most dangerous forms of misalignment you are observing?

MIKE ANDERSON: There is intense pressure from boards and CEOs to accelerate AI adoption. Many organisations are trying to push transformation faster than their operational maturity allows.

Some enterprises already have strong data environments and experimentation cultures. Others are still struggling with foundational readiness. That creates friction across leadership teams because expectations and execution capabilities are often misaligned.

One major concern is the rush toward AI-generated software development without rethinking how future workflows will actually operate.

Many companies are still building applications optimised for humans rather than designing systems for agentic collaboration. Those risks creating fragmented ecosystems filled with disconnected tools, weak integrations, and poor documentation.

Every CIO has seen this problem before. Someone builds an application quickly, it works initially, and later the organisation struggles to maintain or govern it properly.

The other major issue is human fatigue. Machines can generate work at enormous scale and speed, but humans still review many high-risk approvals. If employees suddenly move from reviewing five requests a day to five hundred AI-generated requests, there is a danger they simply begin rubber-stamping decisions.

That creates a completely new category of governance risk.

CIO&Leader: If AI agents increasingly execute workflows, who becomes accountable when hallucinations or security failures occur?



“Many companies are still building applications optimised for humans rather than designing systems for agentic collaboration.”

MIKE ANDERSON: Technology systems have never been flawless, even when operated entirely by humans. Organisations have always dealt with outages, operational failures, and human mistakes.

The same reality applies to AI systems.

Hallucinations can be reduced through better context engineering, stronger workflows, and controlled environments. But the most important safeguard is escalation. If an AI system lacks confidence or encounters ambiguity, humans must enter the process.

Organisations also need realistic expectations around perfection. Humans themselves are not perfectly accurate or unbiased. The real question becomes what level of reliability and governance is acceptable.

Separation of duties remains essential in AI-led systems. An AI agent that writes code should not independently review and deploy that same code into production. Multiple validation layers and approval checkpoints are still necessary, especially for high-risk workflows.

Governance in the AI era is less about removing humans and more about redesigning accountability structures intelligently.

CIO&Leader: Netskope emphasises on adaptive trust and zero trust architectures. In an AI-native world, how do you define trust?

MIKE ANDERSON: Trust today is entirely contextual.

Identity remains foundational. Organisations must know who or what is accessing systems, whether it is a human user or a machine identity. But identity alone is no longer sufficient. Organisations must also understand device posture, browser environments, access locations, application destinations, and the nature of the data being exchanged.

AI introduces another layer of complexity because employees increasingly interact with both enterprise-approved AI systems and personal AI environments. Security teams therefore need visibility not only into where data is being sent, but also what data is being entered into AI systems.

Traditional security models focused on binary “allow” or “block” decisions. Adaptive trust requires a more intelligent approach. Sometimes the right response is coaching users rather than simply denying access.

If an employee attempts to upload sensitive information into a personal AI model, the system should guide them toward the approved enterprise environment instead of only blocking the action.

These same trust principles will eventually apply to AI agents themselves as autonomous systems begin operating using employee credentials, workflows, and enterprise data.

CIO&Leader: How can global enterprises maintain unified security postures amid fragmented regulations across regions like GDPR and India’s DPDP framework?

MIKE ANDERSON: Data sovereignty is becoming one of the defining challenges of enterprise AI adoption.

Organisations must ensure that data generated within specific geographies remains within approved regional infrastructure zones. That requires strong contextual awareness around where data is stored, processed, and transmitted.

For example, an employee in India interacting with enterprise AI systems may need all related processing to remain inside India-based cloud environments.

Security and orchestration systems therefore need to evaluate both the application context and the geographic context of every transaction. If data is moving toward the wrong region, systems should automatically redirect those workflows appropriately.

This level of orchestration is becoming essential as regulatory environments grow increasingly fragmented across global markets.

CIO&Leader: AI infrastructure is expensive. How can CIOs justify

these investments to boards and CFOs?

MIKE ANDERSON: The conversation must begin with outcomes.

AI investments should be treated like long-term strategic infrastructure investments rather than short-term operational expenses. Organisations do not build manufacturing plants expecting profitability on day one. They invest against future business value.

The same principle applies to AI. One of the biggest mistakes organisations make is investing in AI capabilities without clearly defining the business outcomes those investments are expected to generate.

Boards and CFOs think in terms of measurable outcomes, productivity gains, operational efficiencies, and long-term value creation — not experimentation for its own sake.

By 2030, enterprises will likely spend enormous amounts on AI infrastructure, model consumption, and token usage. Someone will need to measure the effectiveness of those investments carefully.

That responsibility will increasingly sit with CIOs.

CIO&Leader: Every vendor today claims to offer AI-driven or zero-trust solutions. Where does marketing diverge from reality?

MIKE ANDERSON: One of the biggest misconceptions is treating zero trust as a single product.

Zero trust is not simply replacing a VPN or deploying one security platform. It is an operational philosophy built around least-privileged access, contextual enforcement, and continuous evaluation.

Many organisations replace legacy systems while still maintaining overly permissive access models underneath. That defeats the purpose entirely. Applying technology to broken processes simply produces the same bad outcomes faster.

Applying technology to broken processes simply produces the same bad outcomes faster.

AI agents will expose those weak-

nesses even more aggressively because they operate at machine speed and identify excessive permissions humans may never notice. Organisations therefore need continuous monitoring, adaptive trust models, and constant reassessment of access behaviours rather than relying on static security assumptions.

CIO&Leader: Looking toward 2030, what do you believe the CIO role will ultimately evolve into?

MIKE ANDERSON: The CIO role will become far broader than technology management. Future CIOs may function as chief innovation officers, chief investment officers, chief integration officers, or even chief risk leaders depending on organisational priorities.

A large part of the role will revolve around human transformation. AI adoption is fundamentally about people, trust, workforce adaptation, and organisational behaviour. Technology leaders must identify employees who are naturally curious about AI and empower them to become catalysts for transformation inside the enterprise.

One major risk I continue to see is organisations rushing to build applications simply because AI-generated development has become easier. The industry must avoid creating fragmented software ecosystems optimised only for humans instead of designing systems built for long-term agentic collaboration.

The vendor ecosystem will also become far more crowded because small teams can now build sophisticated products extremely quickly using AI tools. That makes due diligence, governance, and security validation more important than ever.

Ultimately, the most important trait for future CIOs will be intellectual curiosity. The leaders who succeed will be the ones capable of continuously learning, adapting, and integrating business and technology thinking together. ■

AI without guardrails is just risk at scale

Rohit Shukla, Managing Director, APJ, SolarWinds, discusses how enterprise IT operations are evolving, the growing importance of observability in building resilience, and the path toward more autonomous operations.

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

As enterprises face a growing, complex digital landscape, the organisations are under pressure to quickly move from technology adoption to achieving measurable business outcomes. Amid growing infrastructure complexity, hybrid environments, and rising expectations for always-on services, many organisations still face challenges in visibility, observability, governance, and operational resilience.

In a recent conversation with CIO&Leader, Rohit Shukla, Managing Director, APJ at SolarWinds, an enterprise software company that provides IT management, monitoring, and observability solutions for organisations, discusses how enterprise IT operations are evolving, the growing importance of observability in driving resilience, and the road toward more autonomous operations.

He also delves deeper into SolarWinds' approach for augmenting human decision-making through intelligent automation, the role of its newly launched SW1 Agentic Teammate, and why establishing end-to-end visibility is critical for businesses to modernise their IT environments.

Shukla also reflects on SolarWinds' journey following the Sunburst incident in 2020, the company's continued investments in India, and how it is enabling customers with their AI strategies. Below are edited excerpts from the conversation.

CIO&Leader: AI and agentic AI are gaining strong momentum in enterprises, but many organisations are still unsure about the most impactful use cases. How do you see this market shaping up in India, and what role is

SolarWinds playing in guiding customers through this shift?

ROHIT SHUKLA: There is significant excitement about AI, but also some anxiety. Our latest IT Trends Report, with strong participation from India, shows organisations are approaching AI in two main ways.

The first is improving efficiency and productivity. Enterprises are asking how to automate routine tasks, improve business efficiency, and help employees do more with less.

The second approach is transformational. Organisations are exploring how AI can enable new capabilities, such as entering new markets, creating customer engagement channels, or unlocking entirely new business opportunities.

Meanwhile, security and governance teams are raising key questions: Do we have control? Are there guardrails? Can we trust these systems? Balancing innovation with governance is shaping the current phase of AI adoption.

At SolarWinds, we have already launched agentic AI capabilities on our platform. Our vision is to transition operators into orchestrators, while maintaining a "human in the loop" approach. AI can automate routine tasks, but customers should always retain control and determine their preferred level of autonomy.

CIO&Leader: How are Indian enterprises evolving in their approach to IT operations and observability? What are some of the biggest infrastructure management challenges CIOs face today?



“AI can automate routine tasks, but customers should always retain control and determine their preferred level of autonomy.”

ROHIT SHUKLA: The biggest challenge today is that businesses have virtually no tolerance for downtime.

Boards are less concerned with specific infrastructure components and focus instead on uptime, service levels, and business continuity. CIOs and CTOs are now evaluated more on resilience than on infrastructure management.

Five pillars—governance, compliance, resilience, observability, and autonomy—will define the future of enterprise IT by 2030.

This is where observability is essential. By monitoring, correlating, and analysing the environment, organisations can predict issues and identify root causes more quickly when incidents occur.

The focus has shifted from server and network monitoring to uptime, resilience, and service-level outcomes.

CIO&Leader: How close are enterprises to achieving the vision of autonomous IT operations?

ROHIT SHUKLA: We are still quite far from this goal. The first step is vis-

ibility, as automation is not possible without it. Many organisations lack end-to-end visibility, and some still manage assets in spreadsheets, with separate teams handling networks, servers, and storage independently.

Without a unified view, effective correlation is not possible.

Once organisations achieve visibility, they can analyse incident volumes, alert patterns, and operational bottlenecks. At this stage, AI can serve as a digital assistant, helping teams resolve issues more quickly.

This is a gradual process. Some organisations are in advanced stages, while others are just beginning. Banking and IT services are generally further ahead, while manufacturing, healthcare, and pharma are still in earlier phases.

Every organisation is on its own transformation path.

CIO&Leader: Are we overhyping autonomous IT operations before enterprises are operationally ready?

ROHIT SHUKLA: I would not call it hype, as AI is now a reality. Some organisations will adopt autonomous

capabilities by 2026, while others may not be ready until 2030. The timeline depends on each organisation’s maturity and readiness.

AI will play a major role. Unlike humans, digital assistants provide continuous monitoring and add an extra layer of operational assurance.

The question is not if organisations will adopt AI, but when.

CIO&Leader: The observability market has become increasingly crowded, with vendors specialising in infrastructure monitoring, application performance management, and security operations. How does SolarWinds differentiate itself in such a competitive landscape?

ROHIT SHUKLA: There is no silver bullet in observability. Different vendors solve different problems and have evolved from different starting points. Some vendors built their capabilities around application performance monitoring, focusing on code-level visibility, user experience, and application behavior. SolarWinds, by contrast, established its foundation in infrastructure monitoring, with deep expertise across networks, servers, storage, databases, and hybrid IT environments. While we have since expanded into observability and application monitoring, its heritage remains rooted in infrastructure and IT operations management.

Security vendors also play an important role in the broader IT operations ecosystem by addressing threat detection, compliance, and incident response. As a result, enterprises today are not looking for a single platform that does everything. They are looking for platforms that work together effectively.

For large enterprises, the conversation is rarely about replacing one vendor with another. The real challenge is managing the complexity created by years of department-level technology purchases. Many organisations operate dozens of monitoring, management, and security tools

across different teams, which often leads to fragmented visibility and operational inefficiencies.

Consolidation is now a strategic priority. Customers want to reduce dozens of tools to a few strategic platforms that offer broad visibility and operational functionality without sacrificing functionality.

We see the future as one of coexistence, not replacement. Customers may use different platforms for application performance, security operations, and infrastructure observability. The goal is an integrated ecosystem that correlates data across domains, providing end-to-end visibility and faster problem resolution.

Our strength lies in observability and infrastructure management. Other vendors may have strengths in application performance management or cybersecurity. Ultimately, customers benefit most when these capabilities work together, enabling simplified operations, better resilience, and stronger business outcomes.

CIO&Leader: SolarWinds recently introduced the SW1 Agentic Teammate. How is it enabling more autonomous and self-healing IT environments?

ROHIT SHUKLA: SW1, or SolarWinds One, is the first agent in our broader roadmap, with more agents planned. We have integrated conversational AI into our observability platform, and existing customers can access it by upgrading to the latest version.

The response has been extremely positive. One of the first benefits we noticed was a reduction in support inquiries because customers can now ask the platform questions directly rather than opening support tickets.

We are also seeing customers experiment with entirely new workflows. They are exploring scenarios that previously required specialised expertise and are using the platform in more innovative ways.

“Five pillars—governance, compliance, resilience, observability, and autonomy—will define the future of enterprise IT by 2030.”

SW1 is enhancing efficiency and enabling new operational possibilities.

CIO&Leader: Some industry analysts have warned that vendors sometimes overstate AI's capabilities. Is SW1 truly autonomous, or is it primarily an AI assistant?

ROHIT SHUKLA: Our philosophy is clear: human in the loop. We ensure someone oversees the AI. We have embedded trusted AI principles so customers can see how the AI reached its conclusions, which sources it used, and what information it evaluated.

We see AI as an assistant, not a replacement for humans. The future is people making faster, better, and more informed decisions with AI support. AI is the co-pilot, not the pilot.

CIO&Leader: SolarWinds met significant scrutiny following the Sunburst incident. Have you overcome the trust challenges associated with that event?

ROHIT SHUKLA: The SEC case related to Sunburst was dismissed last year, clearing SolarWinds' name. Notably, the incident mainly affected customers who had not upgraded their platforms despite receiving notifications.

Today, some of India's largest and most heavily regulated banks run their observability environments on SolarWinds. These institutions have rigorous security testing processes and extensive internal controls.

For us, customer confidence is demonstrated by adoption. The continued reliance of highly regulated

organisations on SolarWinds shows our progress in secure-by-design principles and software assurance.

CIO&Leader: Are Indian customers still largely cost-driven buyers, or are they increasingly prioritising resilience and operational visibility?

ROHIT SHUKLA: I don't believe India is a cost-driven market. India is a value-driven market. Customers want the maximum value from their investments. Of course, negotiation is part of the process, but organisations are willing to invest when they clearly understand the business value.

We are not the lowest-cost solution, but we continue to grow because customers recognise the value we provide.

Competing only on price leads to diminishing returns. We prefer to compete based on outcomes.

CIO&Leader: Which industry verticals represent the strongest customer base for SolarWinds in India?

ROHIT SHUKLA: Banking is our largest focus area, followed by IT and IT-enabled services, then pharma, manufacturing, and healthcare. These industries prioritise resilience, observability, and operational continuity, making them natural adopters of our platform.

CIO&Leader: What does SolarWinds' long-term commitment to India look like from an investment and innovation perspective?

ROHIT SHUKLA: India is one of our most strategic global markets. Our largest engineering team, about 500 professionals, is based in India, and our global finance operations are consolidated in Mumbai. We also invest significantly in sales, pre-sales, support, and customer-facing functions.

We see India as critical for both revenue and as a strategic hub for talent and innovation. ■

IT has moved beyond keeping the lights on!

Anjaiah Surgi, CIO, Hindware explains how manufacturing enterprises are scaling AI across plants, dealer ecosystems, forecasting, and operations while balancing cloud infrastructure, data integration, governance, AI literacy, and measurable business outcomes.

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

As manufacturing enterprises accelerate their AI transformation journeys, the focus is rapidly shifting beyond isolated automation projects toward connected intelligence across plants, dealer ecosystems, supply chains, and customer operations. For traditional industries, the challenge is no longer just technology adoption, but building scalable AI-ready infrastructure while balancing operational continuity, data integration, governance, and measurable business outcomes.

In this conversation with CIO&Leaders, Anjaiah Surgi, CIO at Hindware, discusses how the role of the CIO is evolving from infrastructure management to strategic business leadership. He shares how the organisation is leveraging AI, cloud, analytics, computer vision, and predictive intelligence across manufacturing operations, dealer engagement, forecasting, and enterprise decision-making while addressing challenges around AI literacy, security, governance, and sustainable adoption.

CIO&Leader: AI adoption in manufacturing and dealer ecosystems comes with very different operational requirements compared to traditional SaaS environments. How are you balancing low-latency on-ground processing with growing compute requirements?

ANJIAAH SURGI: You cannot continue relying entirely on traditional on-premise infrastructure and static server environments anymore.

That is where cloud and hybrid ecosystems help significantly, whether it is Azure, AWS, or GCP. In our case, we primarily operate on Microsoft

Azure. The workloads are managed dynamically there, which means we do not have to constantly worry about latency or compute scaling.

For example, if I look at our loyalty application ecosystem, we already have more than 60,000 influencers connected through the platform, and the network continues to grow.

If we had depended only on traditional on-premise infrastructure, scaling this environment would have been extremely difficult. Cloud and hybrid environments are what enable us to support this level of scalability and operational flexibility.

CIO&Leader: Legacy manufacturing environments often struggle with fragmented and siloed data. While evaluating technology vendors, how important is integration capability compared to raw AI power?

ANJIAAH SURGI: Data remains the core fulcrum for every decision-making system, including AI.

Whether it is analytics, AI-enabled analytics, or transaction processing, if the underlying data quality is poor, the outcomes will also be poor. It is the classic “garbage in, garbage out” problem.

For example, while creating a sales order, if customer credit limits are not properly updated or integrated across connected systems, it creates operational and financial risks.

When we evaluate vendors or onboard software solutions, scalability, flexibility, and agility become extremely important because technology itself is evolving continuously.

A solution may work extremely well today, but if it cannot evolve alongside security requirements, scalability needs, or architectural changes after one or two years, then it becomes a limitation.

At the same time, integration remains absolutely critical. Our ERP backbone continues to be SAP. Any connected system must integrate properly with the ERP environment. That remains one of the core evaluation parameters for us.

CIO&Leader: Since manufacturing is traditionally not an AI-first industry, how are you addressing the AI literacy gap within the workforce?

ANJIAH SURGI: Any brick-and-mortar industry will face these challenges initially.

What we are doing is conducting continuous education and leadership training sessions across the organisation. We started with leadership teams and gradually expanded across the broader workforce.

Within IT itself, we have created a separate Emerging Technologies department dedicated to AI, ML, IoT, low-code/no-code platforms, and RPA initiatives.

This dedicated function conducts regular enablement sessions across locations throughout India. Over the last year alone, we have conducted more than 20 training and education programs, both virtual and in-person.

The momentum has been very encouraging. Employees are increasingly accepting and actively using AI tools.

Within the Microsoft ecosystem, Copilot access has been provided to employees for tasks such as drafting emails, summarising communication, reviewing quotations, and processing RFPs.

There is still a long journey ahead, but adoption levels are steadily increasing.

We also already have data scientists onboard within the organisation helping drive analytics initiatives across functions.

Some of our business reviews



“AI can automate routine tasks, but customers should always retain control and determine their preferred level of autonomy.”

today happen completely paperless — without PPTs or Excel sheets — directly through BI dashboards.

The next stage we are moving toward is enabling teams to query systems conversationally instead of relying only on preformatted dashboards. As long as the data exists, users should be able to ask questions directly and receive insights dynamically.

CIO&Leader: How are you reshaping dealer and distributor experiences using AI-driven systems?

ANJIAH SURGI: We already operate distributor management systems integrated with our SAP ERP environment.

The next phase is enabling these systems with AI-driven intelligence.

This includes inventory visibility, market intelligence, and predictive recommendations around what dealers should order and when they should place those orders.

Dealers can already access visibility into nearby warehouses and depots from where products can be sourced.

Going forward, we are also building agentic AI-driven workflows where dealers can directly query systems conversationally to check order status, inventory availability, truck movement, and dispatch tracking in real time.

Today many of these processes require navigating traditional transactional systems. We are trying to simplify that experience significantly using conversational AI agents.

Dealers can already place orders through mobile applications. The objective now is to enhance that ecosystem further using AI-driven intelligence.

CIO&Leader: Are these engagement ecosystems aligned with India's DPDP compliance requirements? And are you moving toward hyper-personalisation for dealers and customers?

ANJIAH SURGI: Yes. For all newly developed systems and applications, DPDP compliance is being built in from the beginning.

For legacy systems, we are currently conducting assessments and progressively making them compliant as well.

This includes consent management, data privacy controls, access governance, and related compliance requirements.

On personalisation, yes, we are already moving toward highly personalised engagement models.

Field officers using our applications receive contextual nudges about nearby retailers or business opportunities instead of relying only on static visit plans.

Similarly, dealers and distributors see highly personalised experiences where they only access their own operational data.

We are also increasingly building consent-driven engagement models around anniversaries, milestones, and relationship-oriented communication to strengthen long-term engagement beyond purely transactional interactions.

CIO&Leader: How has the role of the CIO evolved from managing IT assets to becoming a strategic business leader?

ANJIAH SURGI: Traditionally, IT was always viewed as a cost center. That perception is changing rapidly.

Today, boardroom conversations increasingly revolve around how AI improves productivity, drives growth, supports revenue generation, and optimises business performance.

The CIO role is gradually moving closer toward strategic business leadership.

For example, discussions today involve areas such as supply chain intelligence, demand forecasting, production optimisation, and plant efficiency.

In one of our manufacturing facilities, defect detection used to happen manually. About two years ago, we implemented computer vision-based defect detection systems on the production line.

Even with occasional false positives, the accuracy levels are now above 98%, which is significantly better than manual inspection.

Similarly, in demand forecasting,

“Implementing AI alone is not enough. Organisations must continuously measure whether those systems are delivering sustained business value.”

traditional human forecasting accuracy in sanitaryware could sometimes fall to nearly 50–60%.

With AI-enabled forecasting systems, we are now seeing accuracy levels move toward 85–90%.

That directly impacts inventory turnover, inventory carrying costs, manufacturing efficiency, and overall business responsiveness.

Earlier conversations around IT focused on infrastructure upgrades or ERP version upgrades. Today, conversations are centered around plant optimisation, forecast accuracy, regional growth opportunities, and operational intelligence.

These are fundamentally business discussions, and they are already happening.

CIO&Leader: Looking toward 2027, what AI-native business trends do you believe will reshape manufacturing and building material industries?

ANJIAH SURGI: We are already seeing connected plants becoming a major trend.

Many legacy machines that were previously disconnected are now becoming connectable because of evolving technologies.

That enables real-time insights and predictive intelligence for maintenance engineers and production supervisors.

Another major area is defect prediction and quality optimisation.

For example, in ceramic manufacturing, cracks may only become vis-

ible after multiple production stages. Now, using operational data and AI models, we are exploring how to predict and reduce such defects much earlier in the process.

AI models can significantly improve efficiency, reduce defects, and optimise manufacturing quality across production environments..

CIO&Leader: What do you see as the biggest challenges organisations still face while scaling enterprise AI?

ANJIAH SURGI: One of the biggest challenges is balancing innovation with organisational readiness.

Technology is evolving extremely fast — whether it is Claude models, Gemini Enterprise, or other emerging ecosystems. But organisations also need the ability to absorb and operationalise these technologies effectively.

Change management itself becomes a major challenge.

At the same time, security guardrails remain critical. Enterprises must be extremely careful about what technologies they adopt, how enterprise data is protected, and whether governance mechanisms are sufficiently strong.

Another important area is defining proper success metrics and measurable KPIs for AI initiatives.

It is not enough to simply implement AI systems. Organisations must continuously measure whether those systems are delivering sustainable value.

For example, business reviews happening through BI dashboards instead of PPTs only succeed if the organisation sustains that behavioral shift over time.

If momentum drops after a few months, enterprises can easily fall back into siloed operations again.

Without measurable KPIs, sustained adoption becomes extremely difficult — regardless of whether the organisation is a traditional manufacturing enterprise or a high-tech company. ■

NxtGen¹

presents



Building A Cyber Resilient Enterprise

12 -14 June 2026
Hayatt Regency Jaipur Mansarovar

Thank You

To our esteemed CISOs, speakers, industry experts, partners, thank you for making CISO Forum 2026 a remarkable success.

Your trust, insights, and participation fostered meaningful conversations and strengthened a community committed to building a more secure digital future.

PRESENTING PARTNER

NxtGen¹

GOLD PARTNER

ARMIS
from ServiceNow

SILVER PARTNERS



STRATEGIC CASE STUDY PARTNER



ASSOCIATE PARTNER



EXHIBIT PARTNERS



SKILLING PARTNER



EDUCATION PARTNER



BUSINESS INSIGHTS PARTNER



CONCEPTUALIZED AND EXECUTED BY



AMD

presents



ET Edge

27th Annual Conference

CIO&LEADER
 The Agentic Enterprise

PLATFORMS. PEOPLE. POLICY. PROFITS.

co-presented by

NxtGen¹

powered by



co-powered by



freshworks



outsystems

#CIOandLeaderConference

30 **02**
JULY **AUGUST**
2026

FAIRMONT, JAIPUR

For 27 years, ET Edge | CIO&Leader has brought together the minds shaping a future where AI reasons, collaborates, and acts.

Four Days. 200+ India's top CIOs. One Defining Conversation.

WHAT'S IN IT FOR YOU

- ▶ 1:1 meeting with CIOs ▶ Ideas Cafe, Roundtable, Workshop, Booths
- ▶ 2160+ Networking minutes ▶ CIO Samman & NEXT100 Awards ▶ Musical evening with Rajasthani folk group

EMINENT PERSONALITIES

**Toby Walsh**

Laureate Fellow &
Scientia Professor of AI,
University of
New South Wales

**Rana Ashutosh Kumar Singh**

MD (International
Banking, Global Markets
& Technology), SBI

**Madan Sunder Das**

Monk, ISKON, Founder,
Spiritual Guide and
Community Leader,
EVOLVE Pune

**Jim Sarbh**

Indian film and
theatre actor

**Abhijeet Sawant**

Indian Singer and the
winner of the first
season of Indian Idol

Be Where India's Tech Leaders Meet.



For Sponsorship, Write to

Naveen Singh

c-naveen.singh@timesgroup.com
9901300772

Hafeez Shaikh

hafeez.shaikh@timesgroup.com
9833103611

Aanchal Gupta

aanchal.gupta1@timesgroup.com
9651841119

Supriya Sahoo

supriya.sahoo@timesgroup.com
8095056886