



CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

A 9.9 GROUP PUBLICATION
cioandleader.com  cioandleader  cioandleader/

THE AGENTIC ENTERPRISE When AI acts, who stays in control?



**AJAY KUMAR
AJMERA**
Group CIO
Rockman Industries



**DEBANJAN
BANERJEE**
Global Service Management Director
Pernod Ricard India



**GAURAV
DUGGAL**
SVP- IT & Security
Jio Platforms



**INDRADYUMNA
DUTTA**
Group CDO & GCC Head
Jindal Steel and Power



**RITESH
KUMAR**
Assistant Vice President - IT
EXL



**SUPRIYA
KAUL**
IT Regional CIO (APAC)
CNH Industrial



CIO&LEADER



studioTalks



+ CIO&LEADER studioTalks

CIO&LEADER STUDIOTALKS— WHERE TECHNOLOGY MEETS THE SPOTLIGHT!

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

WHY JOIN STUDIOTALKS?

Engage in powerful conversations that shape the future of enterprise IT.

Share your expertise in a high-impact, TV-style format.

Be featured among India’s top technology leaders.

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

SECURE YOUR SPOT NOW!

For more information
Jatinder Singh

Chief Editor, Enterprise Tech Publications
ET Edge - Times Group

jatinder.singh1@timesgroup.com, +91 9718154231

For Business Proposal
Hafeez Shaikh

Assistant Director - Projects
ET Edge - Times Goup

hafeez.shaikh@timesgroup.com, +91 9833103611

Follow us: @CIOandLeader



Continuity, with greater scale!

Starting March 1, 2026, CIO&Leader has moved from 9.9 Group to ET Edge. This transition also includes the CISO Forum and ITNEXT brands. While it reflects a change in ownership, more importantly, it signals the start of a new phase—one defined by greater scale, expanded reach, and a stronger ecosystem to drive future growth.

What remains unchanged is our core strength. The brand, the editorial philosophy, and the team you know and trust remain unchanged. The relationships we have built with you through consistent, practitioner-led engagement remain intact.

What changes is the stage! As part of ET Edge, an initiative of The Times Group, CIO&Leader now operates within one of India's most influential media and business ecosystems. This is not just about expanded reach. It places enterprise technology conversations within the broader narrative of business strategy, economic direction, and industry transformation.

In many ways, my own journey with this platform reflects that same continuity. I first joined in 2009, during the early days of CIO&Leader (then The CTO Forum) and IT Next, and spent nearly 3 years helping shape their foundations. After some time away, I returned post-COVID as Editor, reconnecting with a platform that had evolved

alongside its community. Today, as Chief Editor, that journey comes full circle, rooted in the same purpose, now with a wider canvas.

That sense of continuity extends beyond individuals to the platform itself. Over the years, CIO&Leader, along with NEXT100 and CISO Forum, has grown alongside the enterprise technology ecosystem.

Over the past 25 years, many of you have been part of this journey in different roles as participants, contributors, jurors, and winners. That continuity speaks to both the strength of the platform and the way this community has evolved together.

Looking ahead, our priority remains consistency in editorial quality, depth, and engagement.

At the same time, you will see a broader footprint, not through an intent shift but through expanded reach and relevance. The objective is clear: to ensure that the insights, experiences, and leadership emerging from this community are amplified at the scale they deserve.

It is also important to acknowledge those who have shaped this journey. I thank Vikas Gupta, Co-founder of 9.9 Group and former Editorial Director, for his vision and leadership in building these platforms. As he moves into retirement, his contribution remains foundational to what CIO&Leader represents today.

To our CIO and CISO community, thank you for your

“As part of ET Edge, an initiative of The Times Group, CIO&Leader now operates within one of India’s most influential media and business ecosystems, amplifying its reach and impact among business and technology leaders.”



continued trust and engagement. As we step into this next phase, there is both excitement for what lies ahead and a sense of nostalgia for the journey so far. With your continued support, we look forward to taking this platform to new heights. ■

Jatinder Singh
Editor
jatinder.singh1@timesgroup.com



COVER STORY

10-16

The Agentic Enterprise: When AI acts, who stays in control?



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

06

Inside the AI Impact Summit 2026 and what it means for IT leaders



INSIGHT

17-18

The modern vehicle is a computing system on wheels



19-21

Why enterprises must move from experiments to outcomes



TECH TALK

22-24

How CIO accountability is being rewritten for...

VIRAJ DESHPANDE



25-26

CIOs are now accountable for outcomes, not infrastructure

ASHISH THAKUR



27-30

You can't bet your business on AI you can't verify

TIRTHANKAR LAHIRI



31-34

Why CIOs must rethink their data stack in 2026?

HEMANT TIWARI



35-37

The new CIO scorecard: Revenue, resilience, and responsible AI

PREMKUMAR BALASUBRAMANIAN

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**

Principal Correspondent & Editorial Coordinator –

CIO&Leader: **Musharrat Shahin**

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**





**Kuldeep Koul
Appointed CDO at
NAMTECH**

Kuldeep Koul joins NAMTECH as Chief Digital Officer to lead its digital strategy and transformation.



**Amit Waghmare
Takes Over as
CIO at Bennett
Coleman**

Amit Waghmare has been appointed CIO at Bennett Coleman to drive enterprise IT modernization.



**Vijay Narayanan
Appointed Head
of Innovation &
AI at Kotak
Mahindra Bank**

Vijay Narayanan joins Kotak Mahindra Bank to lead innovation and AI initiatives.



**Somak Shome
Joins JAC OLIVOL
as Head of
Enterprise IT**

Somak Shome joins JAC OLIVOL as Head of Enterprise IT strategy and Transformation.



**Gyan Pandey
Appointed
Executive
President & CDIO
at Polycab India**

Gyan Pandey joins Polycab India as Executive President & CDIO to drive digital and data strategy.



**Vinay Kumar
Appointed Head of
IT at Tim Hortons
Middle East**

Vinay Kumar joins Tim Hortons Middle East to lead IT and digital transformation.



**Pramod Adiddam
Appointed CTO at
Myntra**

Pramod Adiddam joins Myntra as CTO to lead platform, data, and AI strategy.



**Kallol Basu
Appointed
Divisional CIO
at ITC**

Kallol Basu joins ITC to drive corporate IT transformation.



**Rahul Dayal
Appointed CTO at
SBI Mutual Fund**

Rahul Dayal joins SBI Mutual Fund as CTO to scale digital and core platforms.



**Brijesh Mishra
Appointed VP –
Sales & Revenue at
SecHard**

Brijesh Mishra joins SecHard to lead sales and revenue across Asia.



Purvi Shah
Appointed Head of IT at Ajmera Realty

Purvi Shah joins Ajmera Realty to lead IT modernization and digital platforms.



Kartikeya Pandey
Appointed AGM – IT at Kajaria Ceramics

Kartikeya Pandey joins Kajaria Ceramics to strengthen IT operations and drive transformation.



Gaurang Vora
Appointed SVP & Head of Technology at GlobalPay WSF_x

Gaurang Vora joins GlobalPay WSF_x to lead technology modernization.



A. N. Srinivasan
Promoted to SVP – IT at SRF

A. N. Srinivasan has been elevated to SVP – IT at SRF to lead ERP modernization.



Ravi Kumar Pangal
Appointed CIO at IndusInd Bank

Ravi Kumar Pangal joins IndusInd Bank as CIO to modernize core banking technology.



Rahul Agarwal
Appointed VP – IT at MakeMyTrip

Rahul Agarwal joins MakeMyTrip to lead enterprise IT and digital platforms.



Taposh Saha
Appointed SVP – IT at HDFC Bank

Taposh Saha becomes SVP – IT at HDFC Bank to drive digital banking initiatives.



Upkar Singh
Appointed Head of IT at Paradeep Phosphates

Upkar Singh joins Paradeep Phosphates to lead digital transformation and IT strategy.



Krishan Bhardwaj
Appointed Group IT Head at MGM Healthcare

Krishan Bhardwaj joins MGM Healthcare to lead group-wide IT modernization.



Sandeep K Vasudevan
Appointed Head – IT Infrastructure at Wipro Enterprises

Sandeep K Vasudevan joins Wipro Enterprises to lead infrastructure and cloud strategy.

AI Impact Summit 2026 signals next phase of enterprise AI adoption

Global scale, local execution, and the infrastructure push driving enterprise AI

By **Musharrat Shahin** | editor@cioandleader.com

The India AI Impact Summit 2026, held at Bharat Mandapam from February 16 to 20, was among the most ambitious artificial intelligence gatherings in the Global South. The five-day event brought together global policy-makers, industry leaders, research innovators, and enterprise technologists to discuss AI's future, its opportunities, risks, and the investments shaping its development.

Global participation

Prime Minister Narendra Modi inaugurated the summit, positioning India as a global hub

for AI innovation focused on public benefit and inclusive growth. He noted that India's extensive digital ecosystem and growing talent pool uniquely position the country to influence global AI deployment.

The Summit drew delegations from more than 100 countries and included over 20 heads of state and 60 ministers, underscoring its diplomatic technological significance. Notable attendees ranged from French President Emmanuel Macron to leaders of major AI firms.

Senior executives from leading global tech companies, including Sundar Pichai (Google) and Sam Altman (OpenAI), participated in panels and plenary sessions exploring AI's commercial opportunities, governance challenges, and societal impacts. Strategic commitments were a central theme. One of the biggest announcements was a collaboration between OpenAI and Tata Consultancy Services (TCS) to build 100 MW of dedicated AI compute infrastructure in India to support large-scale AI workloads and bolster the country's global competitiveness.

Reliance Industries and its telecom subsidiary Jio pledged approximately \$110 billion to AI and data center infrastructure over the coming years. This investment aims to address local compute shortages and expand capacity for enterprise and public sector adoption. The AI impact summit highlighted that AI is emerging as a national priority, with governments and enterprises accelerating investments in infra, policy and innovation ecosystems. ■



India AI Impact Summit was among the most ambitious AI gatherings in Global South

India leads global GenAI returns as 71% of firms report positive ROI

Snowflake and Omdia data show 71% of Indian firms are seeing positive returns from generative AI initiatives

By **CIO&Leader** | editor@cioandleader.com

Indian enterprises are moving past the experimental phase of artificial intelligence to achieve measurable financial gains. A new global research report from Snowflake and Omdia, “The ROI of Gen AI and Agents,” indicates that 71% of Indian organisations now report a positive return on investment from generative AI. This figure stands ten percentage points higher than the global average of 61%.

The survey, which gathered data from 2,050 technology leaders across ten countries, places India among the top three global adopters alongside the United States and Germany. Currently, 42% of Indian companies utilise generative AI across multiple departments.

High adoption in technical workflows

The data shows that Indian firms apply these tools heavily within technical functions. While global adoption in software development sits at 48%, Indian firms report a 65% adoption rate. This trend extends to several core operations:

- **IT Operations:** 76% in India against 61% globally.
- **Data Analytics:** 75% in India against 57% globally.

Indian enterprises plan to allocate 28% of their total technology budgets to GenAI

- **Cybersecurity:** 69% in India against 52% globally.

These figures suggest that Indian businesses are prioritising under-the-hood improvements to their digital infrastructure. Organisations by automating routine engineering tasks, are aiming to shorten development cycles and improve system reliability.

Measurable business outcomes

According to the report, 94% of Indian respondents noted gains in operational capacity, while 86% reported direct improvements in financial performance.

Specific use cases in software engineering illustrate this trend. In India, 77% of organizations use AI for code generation, and 76% use it for debugging and reviews. In contrast, global averages for these tasks are 64% and 65%, respectively.

Projected spending and agentic AI

India is also expected to lead global spending on these technologies over the next year. Indian enterprises plan to allocate 28% of their total technology budgets to generative AI initiatives. This is the highest planned allocation among all surveyed nations and exceeds the global average of 22%.

The next phase involves agentic AI systems capable of executing complex tasks with minimal human intervention. While 56% of global organisations plan to implement these agents, 66% of Indian firms have already started or intend to start within the next 12 months. ■

HPE and NVIDIA unveils next-gen supercomputing systems

HPE and NVIDIA expand AI portfolio with Vera Rubin architecture, liquid-cooled supercomputers, and air-gapped security

By **CIO&Leader** | editor@cioandleader.com

Hewlett Packard Enterprise (HPE) announced a broad expansion of its NVIDIA AI Computing by HPE portfolio. The updates introduce new hardware based on the NVIDIA Vera Rubin architecture and specialised security features for sovereign and regulated environments.

The new HPE Cray Supercomputing GX240 compute blade supports up to 16 Vera CPUs per blade. This system uses 100% fanless direct liquid cooling to manage heat while maintaining high hardware density.

A single rack of the GX240 can scale to 640 Vera CPUs and 56,320 ARM-compatible cores. This infrastructure is built for the most intensive tasks, such as training trillion-parameter large language models (LLMs) and running agentic AI, autonomous agents that perform complex reasoning and task execution.

Scaling the AI factory

The HPE AI Factory portfolio now includes the NVIDIA Vera Rubin NVL72 rack-scale system. This configuration integrates 36 Vera CPUs and 72 Rubin GPUs. Compared to previous Blackwell-based systems, this architecture offers up to a 10X reduction in costs for inference tasks and requires four times fewer GPUs to train Mixture-of-Experts (MoE) models.

HPE also introduced the HPE Compute XD700, an AI server based on the NVIDIA HGX Rubin NVL8.

This system provides double the GPU density of previous generations, supporting up to 128 GPUs per rack.

For research institutions and sovereign entities, new blades have been introduced for the Cray Supercomputing GX5000 platform. These include:

- **Vera CPU Compute Blade**
- **Rubin NVL72 System**
- **Liquid Cooling**

Security for sovereign and private clouds

In terms of meeting the needs of regulated industries like defence and finance, HPE Private Cloud AI now offers an air-gapped configuration. This ensures that sensitive data remains isolated from external internet connections. The hardware integrates specialised data processing units (DPUs) and high-speed Ethernet networking. These components offload tasks from the main processors, which improves overall system speed. On the security front, the systems include post-quantum cryptography and FIPS 140-3 Level 3 certification to protect sensitive data against emerging digital threats.

The private cloud system can now scale up to 128 GPUs through new network expansion racks. It also integrates NVIDIA Mission Control software, which automates workload orchestration and system recovery. ■

A digital coworking room without humans? How does this work

AI is evolving from a helper to an autonomous coworker that can plan and complete tasks on its own

By **CIO&Leader** | editor@cioandleader.com

The idea of a digital coworking room without humans once sounded like science fiction. Today, it's becoming reality. With the arrival of tools like Claude Cowork by AI giant Anthropic, AI is no longer just assisting work; it is starting to do the work. Investors quickly reacted to the disruption Claude Cowork could bring to traditional IT services. This led to a broad sell-off in global tech stocks.

From chatbots to autonomous digital coworkers

Claude Cowork is an open-source plugin designed to help enterprises build autonomous workflows. Unlike traditional chatbots that simply respond to prompts, this tool can be given controlled access to folders and files and then plan, execute, and complete multi-step tasks on its own. From organising documents and drafting reports to creating spreadsheets from scattered data, AI is now acting more like a digital employee than a virtual assistant.

The conversation around AI has clearly moved from AI helping humans work faster to AI being capable of doing the work itself.

IT stocks feel the heat

The impact wasn't limited to tech discussions; it quickly reached the stock markets. Investors reacted sharply to the disruption this kind of AI could bring to traditional IT services, an industry long dependent on human billable hours.

Globally, tech and software stocks saw a sell-off. In India, the NIFTY IT index dropped around 6–8%, wiping out billions in market value from companies like Infosys, TCS, and Wipro. The concern was clear: if AI can autonomously handle tasks, demand for conventional service-based IT contracts could decline.

So, where do humans fit in?

Despite these fears, humans are far from irrelevant. AI may execute tasks faster, but humans still play a critical role in setting direction, making strategic decisions, and understanding context. Most importantly, humans bring judgement and ethics to areas where AI still falls short.

AI can analyse data and suggest actions, but it cannot fully understand fairness, social impact, or moral responsibility. Deciding what is acceptable, setting boundaries, and taking accountability must remain human responsibilities.

The real challenge

As autonomous AI grows, ethical and responsible use becomes the biggest concern. Humans must ensure AI systems are fair, transparent, and aligned with societal values. AI may run on algorithms, but responsibility cannot be automated. Claude Cowork doesn't signal the end of human work; it signals a transformation. The future belongs to organizations that learn how to combine autonomous AI with human judgement, creativity, and ethics. ■

THE AGENTIC ENTERPRISE

When AI acts, who stays in control?

By **Musharrat Shahin** | editor@cioandleader.com

At 2:17 a.m. on a late January night, long after the network team had signed off, a closed-loop AI agent inside a telecom network noticed something unusual. There was a subtle but consistent spike in latency across multiple nodes. It did not escalate the issue or wait for intervention. Instead, it analyzed the anomaly in real time, rerouted traffic, adjusted network parameters, and triggered corrective actions, restoring stability before users noticed any impact.

By morning, everything was running smoothly. No late-night calls. No war rooms. No frantic email threads. Just business as usual. Incident resolution times dropped by 30 to 40 percent.

But here is the part that makes most leaders pause: no human stepped in. No one approved those actions at that moment.

A few weeks later, at an industrial company, a dynamic pricing system told a different story. It had performed well in testing, but early in production it ran into a messy data issue. Nothing major, but enough to put margins at risk.

The system did not fail, but trust did. The project was paused, not because the technology broke, but because the organization was not comfortable letting it run without tighter control.

For most CIOs and tech leaders, this situation feels familiar. For years, AI helped analyze data, predict outcomes, and optimize decisions, but humans still made the final call. That line is starting to blur.

Today's AI agents do not just recommend, they act. They are appearing inside ERP systems, supply chains, and factory floors. In manufacturing, they monitor equipment and trigger fixes. In enterprises, they are beginning to make decisions that directly impact operations.

And that is where it gets real. The question is no longer can AI do this? The real questions are:

- Are we ready to live with what it does?
- How much decision-making power are we willing to hand over?
- What level of risk feels acceptable when the system acts independently?
- And when something goes wrong, who is accountable?

Technology is moving fast. The harder part is deciding how much control to give up and how much responsibility to take back.

Why scaling is failing

Industry signals indicate enterprises are entering the early stages of the agentic AI era, but adoption remains deliberate. According to McKinsey & Company's latest global AI survey, over 70 percent of organizations now use AI in at least one business function, yet only a small fraction have scaled it enterprise-wide.

The gap is revealing. Data from the CIO&Leader CIO Priorities Survey reinforces this trend: more than 80 percent of CIOs rank AI as a top priority, yet it still trails cybersecurity. Enterprises recognize AI's potential but remain cautious about operational and financial risks.

“We do not allow any AI agent to independently create a material impact on the balance sheet or take actions that could cause significant reputational harm.”

Gaurav Duggal
SVP – IT & Security,
Jio Platforms



“Our operating principle is simple: where errors carry high or irreversible consequences, humans retain final accountability.”

Ajay Kumar Ajmera
Group CIO,
Rockman Industries



This caution reflects fundamental questions: How do you govern systems that can act independently? Where should autonomy stop? Who is accountable when AI decisions carry risk?

Insights from industry leaders illustrate how this transition is unfolding. At CNH Industrial, the approach is deliberate and controlled. Supriya Kaul, IT regional CIO for APAC, explains: “We are running pilots on AI. We are not putting them in full-blown production yet.” AI agents handle infrastructure monitoring—tracking server utilization, RAM consumption, and system alerts—tasks that once required constant manual oversight. Nearly 80 percent of monitoring activities are automated, freeing IT teams to focus on strategic priorities.

Beyond operations, AI-driven analytics inform business decisions. By analyzing regional demand patterns and product performance, the company has identified underserved markets, contributing to a 15–20 percent uplift in sales. Boundaries remain clear. “We are not letting AI take legal or compliance decisions,” Kaul emphasizes. Customer interactions, regulatory responsibilities, and sensitive decisions remain under human control. The goal is confidence, not speed—building trust in AI in controlled environments before expanding its authority.

Debanjan Banerjee, Global Service Management Director at Pernod Ricard India, echoes this pragmatic approach. His team leverages AI solutions embedded within enterprise-grade platforms like Workday and ServiceNow, ensuring continuous product evolution and quality outcomes without heavy internal overhead. Current use cases include IT service management, HR processes, transversal knowledge search, and cybersecurity.

Enterprises with a history of AI deployments have seen significant impact in sales, marketing, and operations. AI has optimized marketing spend, improved demand fulfillment, and enhanced inventory management, directly influencing both P&L and balance sheets. However, high-stakes financial or reputational decisions have not yet been delegated to autonomous AI agents. Organizations are preparing for this next stage by building governance frameworks, controls, and risk thresholds.

“Currently, AI agents focus on low-risk, efficiency-driven tasks where productivity gains are clear and compliance is maintained. Enterprise deployment requires careful evaluation of general-purpose versus domain-specific models, ROI, and sustainability,” Banerjee adds.

Looking ahead, AI agents will increasingly manage workflow orchestration, routine approvals, operational optimization, ticket routing, expense approvals within thresholds, and personalized onboarding and coaching. Strategic, ethical, and high-stakes decisions will remain human-led. “In three years, AI agents will likely manage sub-processes autonomously with frameworks for explainability, auditability, and risk calibration, while humans retain oversight on strategic decisions,” he says.

Where AI is working

Across industries, enterprises are moving beyond experimentation with AI agents, demonstrating how the technology can drive operational efficiency and strategic advantage. Sectors such as telecom, manufacturing, logistics, and retail are increasingly embedding AI into core workflows—handling everything from network monitoring and predictive maintenance to supply chain optimization and

AI Agent adoption journey

Company	Stage of AI agent deployment	Key focus areas	Level of autonomy	Impact and notes
Jio Platforms	Production-scale	Network operations, customer support, fraud detection, IT service management	Tiered: fully autonomous (low-risk tasks), human-in-the-loop (moderate-risk tasks), human-controlled (high-risk tasks)	30–40% reduction in incident resolution; improved service reliability and operational efficiency
Rockman Industries	Pilot and selective deployment	Manufacturing processes, predictive maintenance, supply chain optimization	Human-monitored; initial pilots in low-risk operational scenarios	Early improvements in machine performance monitoring, predictive maintenance; risk-managed deployment
EXL	Controlled pilot	Data analytics, workflow automation, ITSM, HR processes, cybersecurity	Human oversight mandatory; outputs monitored and validated	Productivity gains, automation of repetitive tasks; AI viewed as productivity accelerator, not decision-maker
Jindal Steel & Power	Production and integration	Heavy manufacturing, blast furnace optimization, shop-floor safety, predictive maintenance	AI supports decisions but humans retain strategic oversight	Improved efficiency, reduced downtime, enhanced safety; predictive and prescriptive intelligence enabled
CNH Industrial	Pilot phase	IT operations monitoring, analytics	AI agents handle monitoring; strategic, regulatory, and compliance decisions human-controlled	Nearly 80% of monitoring automated; strategic decisions remain human
Pernod Ricard India	Enterprise-grade platform deployment	ITSM, HR processes, transversal knowledge search, cybersecurity	Limited autonomy through embedded AI platforms	Measurable productivity gains; controlled AI adoption to minimize internal overhead

“Autonomy is deliberately limited: outputs are monitored and validated, and human oversight remains central. AI augments productivity without compromising control.”

Ritesh Kumar
Assistant Vice
President – IT, EXL



personalized customer engagement. Companies like Jio Platforms and Rockman Industries exemplify this shift, showing how disciplined AI adoption can deliver measurable outcomes while maintaining control and accountability.

Gaurav Duggal, SVP – IT & Security at Jio, notes that AI agents are no longer pilots—they are embedded across core operations, driving automation and real-time decision-making. “We are well beyond experimentation. AI agents are operating in live production across several core functions,” he says.

AI is delivering tangible results in customer support, network monitoring, fraud detection, marketing personalization, IT service management, and supply chain forecasting. A mature example is the closed-loop network operations agent, which ingests telemetry, detects anomalies, identifies root causes, and autonomously takes corrective action. Properly implemented, it can cut incident resolution times by 30–40 percent in select clusters while boosting reliability and operational efficiency.

- Scaling AI safely requires disciplined control. At Jio, decision-making is tiered by risk:
- Fully autonomous: low-risk tasks like network tuning or Level-1 issue resolution
- Human-in-the-loop: moderate-impact decisions like pricing adjustments or fraud checks
- Human-controlled: high-stakes choices affecting compliance, capital, or reputation

Ajay Kumar Ajmera, group CIO at Rockman Industries, emphasizes a similar approach in manufacturing. Production processes, supply chains, and quality control must work in sync, so AI integration requires careful alignment with existing systems.

The first step is a strong digital foundation. Operational data—from

shop-floor machines to ERP systems—must be structured, reliable, and accessible. Once in place, AI can analyze machine performance, detect potential disruptions, enable predictive maintenance, and reduce unplanned downtime. Beyond the shop floor, AI can optimize supply chain planning, inventory, and scheduling, improving throughput and resource utilization.

Ajmera advises starting with well-defined scenarios where outcomes are measurable and risks are manageable. This disciplined approach allows enterprises to scale AI effectively, delivering value without compromising control or accountability.

The governance gap

AI adoption in enterprises is moving from experimentation to practical deployment, but success depends on governance and striking the right balance between automation and human oversight. Organizations must determine which tasks AI can handle independently and which require human judgment, ensuring efficiency without compromising control, accountability, or compliance. This balancing act is not just internal, governments and regulators are also shaping the ecosystem. Frameworks such as the EU’s AI Act, evolving national AI strategies, and industry-specific mandates (for example in financial services and healthcare) are pushing enterprises to align innovation with legal and ethical standards before scaling agentic systems broadly.

Ajay Kumar Ajmera emphasizes this in manufacturing. Routine tasks, such as monitoring equipment or analyzing production data, can be automated, but strategic decisions—on capital allocation, supplier selection, and major operational changes—remain under human control. The goal is a collaborative ecosystem where AI enhances

“AI, IoT, and predictive analytics are improving efficiency, reducing downtime, and enhancing safety across industrial operations.”

Indradyumna Dutta
CDIO,
Jindal Steel & Power



“We are running pilots on AI. We are not putting them in full-blown production yet. We are building confidence in controlled environments before expanding its authority.”

Supriya Kaul
IT regional CIO for APAC,
CNH Industrial



efficiency while humans provide oversight, strategic direction, and compliance assurance.

Ritesh Kumar, assistant vice president – IT at EXL, highlights a parallel approach in analytics. AI agents developed through the company’s AI Centre of Excellence automate repetitive tasks like documentation, data preparation, and workflow coordination. Autonomy is deliberately limited. Outputs are monitored and validated, and human oversight remains central. Strong governance, including access controls, continuous monitoring, and override capabilities, ensures AI augments productivity without undermining accountability.

Indradyumna Dutta, CDIO at Jindal Steel & Power Ltd., illustrates how AI, IoT, and predictive analytics transform heavy manufacturing operations. At Jindal Steel, AI supports processes across the entire steel value chain. Predictive models optimize blast furnace productivity, reducing fuel consumption and improving output efficiency. Computer vision systems monitor shop-floor activity in real time, ensuring safety compliance and detecting hazards before incidents occur. IoT sensors and predictive maintenance algorithms identify potential equipment failures before they happen. These capabilities improve asset reliability, reduce downtime, and enhance safety across industrial operations.

The company’s digital architecture integrates operational technology with enterprise IT systems, enabling real-time data flow across production, logistics, finance, and supply chain operations. This integration allows leadership to move beyond monitoring toward predictive and prescriptive operational intelligence while maintaining regulatory compliance and alignment with government-mandated safety and reporting standards.

Beyond manufacturing, financial firms provide blueprints for harmonizing AI independence with stringent oversight. Institutions like JPMorgan Chase and HSBC layer AI governance into their risk systems, fueling massive deployments alongside human reviews, compliance audits, and multidisciplinary councils. At JPMorgan, AI controls mesh seamlessly with enterprise-wide risk protocols, expanding as adoption grows; HSBC, meanwhile, leverages AI pipelines to forecast compliance hurdles and fairness issues across global markets. This integration powers breakthroughs in fraud prevention and customer experience, proving governance can be operational bedrock rather than a bolt-on.

Around the world, companies are rising to regulatory challenges with dedicated ethics committees and responsible AI systems. Microsoft’s guiding principles, Roche’s structured processes, and Walmart’s AI tools stand out—they embed transparency, equity, and responsibility into tech creation, while tracking ethical and sustainability metrics in sprawling supply chains. These efforts mark AI governance’s evolution into a pillar of corporate risk, reporting, and executive scrutiny, far from mere tech silos.

In essence, these models reveal the formula for trustworthy, high-impact AI: disciplined controls, defined responsibilities, and regulatory sync, all sustaining autonomy’s productivity edge.

Looking ahead

The journey of agentic AI is no longer theoretical; it is happening in real time across industries. From telecom networks that resolve incidents autonomously to factories and analytics firms where AI augments human decision-making, the technology is proving its value. Yet the lesson is clear: capability alone

“We have not delegated high-stakes decisions to AI, instead building governance, controls, and risk thresholds to enable responsible adoption.”

Debanjan Banerjee
Global Service Management
Director, Pernod Ricard India



is not enough. Success depends on governance, disciplined control, and a careful balance between autonomy and oversight.

For CIOs and technology leaders, the question is not whether AI can act independently, but whether organizations are ready to trust it, manage risk, and retain accountability. The companies leading the way are those that start with manageable risks, build strong digital foundations, and design AI to complement human expertise rather than replace it. The strategy is clear: start where risks are manageable. Early deployments of agentic AI focus on internal IT operations, where AI agents monitor systems, automate routine tasks, and deliver measurable value without compromising trust or compliance. With the right guardrails, agentic AI can extend traditional AI frameworks into areas where autonomy amplifies efficiency while maintaining accountability.

In this balance between machine intelligence and human judgment lies the true promise of the agentic enterprise: systems that act decisively when needed, while humans provide guidance, context, and oversight. With the right guardrails, such scenarios are not outliers, they represent the future of enterprise operations.

The journey of agentic AI is no longer theoretical; it is unfolding in real time across industries. From telecom networks that resolve incidents autonomously to factories and analytics firms where AI augments human decision-making, the technology is already demonstrating tangible value. Yet one message is becoming increasingly clear: capability alone is not enough. Success will depend on governance, disciplined control, and a careful balance between autonomy and oversight.

For CIOs and technology leaders, the real question is no longer whether AI can act independently, but whether organizations are ready to trust it responsibly, manage the risks it introduces, and retain accountability for its outcomes. The leaders in this space are not rushing blindly into scale. They are deliberate. They begin with manageable risks, invest in strong digital foundations, and design AI systems that complement human expertise rather than replace it.

This is why many enterprises are starting within controlled environments, particularly in internal IT operations, where AI agents can monitor systems, automate routine tasks, and deliver measurable value without compromising trust or compliance. These early deployments are not just about efficiency; they are about building confidence, refining guardrails, and understanding failure scenarios before expanding AI's authority.

But as adoption scales, the stakes will rise. Governance will move from being a technical consideration to a board-level priority. Organizations will need to embed auditability, security, and cost discipline directly into their AI architectures, not as afterthoughts but as foundational principles. The ability to explain decisions, trace actions, and intervene when necessary will define whether enterprises can scale AI sustainably.

Ultimately, the promise of the agentic enterprise lies in balance. Machines will act faster and with greater precision, but human judgment will remain essential to provide context, direction, and ethical oversight. Enterprises that recognize this balance and build for it intentionally will not only unlock efficiency but also earn the trust required to lead in an AI-driven future. ■



The modern vehicle is a computing system on wheels

Rajiv Pandey, CTO and VP, Tata Motors, on how software, data, and AI are driving modern business

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

For much of the past century, the car was defined by mechanical engineering. Power, durability, and manufacturing scale determined success. That hierarchy is now shifting. Software, data, and artificial intelligence are beginning to shape how vehicles are designed, built, and operated.

Tata Motors, says that this transition has unfolded steadily rather than suddenly. The company says that it has been systematically collecting data across vehicles, manufacturing plants, dealers, service networks, and parts logistics since two decades. What started as operational record keeping has since evolved

into a system that increasingly influences how decisions are made.

“Data started as something we observed,” says Rajiv Pandey, chief technology officer and vice-president of Tata Motors. “Today, it is something that drives the business.”

That shift, he argues, mirrors a broader change across the automotive industry. “The modern vehicle is no longer just a mechanical product. It is a computing system on wheels, making decisions in real time.”

An industry redefined by software

The automotive sector is undergoing one of its most consequential transitions since the advent of mass production. Electrification, emissions regulation, and changing consumer expectations are forcing manufacturers to compete on intelligence as much as on hardware.

Industry estimates suggest that software and digital services, which currently account for a modest share of automotive revenues, could represent a majority of industry value within the next decade.

From data to operational leverage

At Tata Motors, AI has moved beyond experimentation roughly six years ago. The focus shifted from isolated proofs of concept to embedding analytics directly into core workflows.

“We realised AI could not sit on the side as a pilot,” Pandey says. “It had to be part of how logistics, supply chains, and manufacturing actually work.”

Logistics provided an early demonstration of impact. By digitising planning and applying predictive algorithms, the company reduced the time required to move parts from warehouses to assembly lines from nearly twelve days to about



Rajiv Pandey
CTO and VP,
Tata Motors

one and a half. The improvement was not just of speed. Faster parts movement helped stabilise production schedules and reduce inventory risk, particularly during periods of supply-chain disruption.

Similar approaches were applied to parts planning and distribution. These efforts improved availability and contributed to a marked expansion of the parts business, while also lowering the likelihood of production stoppages.

Pandey discussed these changes recently while speaking at the Oracle AI World Tour in Mumbai, placing Tata Motors’ experience within a wider industry context.

Intelligence inside the vehicle

The application of AI is equally visible in Tata Motors’ products. Modern vehicles increasingly rely on software to manage safety, performance, and energy consumption. Multiple systems operate simultaneously, taking real-time decisions on braking, airbags, battery health, and range optimisation.

“In today’s vehicles, multiple AI systems are continuously running in the background,” Pandey explains. “They are making deci-

sions that directly affect safety and efficiency.”

Electric vehicles intensify this trend. Managing charging behaviour, predicting usable range, and maintaining battery performance under varied conditions require continuous analysis of data. In this environment, AI becomes less an optional enhancement and more an operational necessity.

Scaling these capabilities has required significant investment in digital infrastructure. Tata Motors has built platforms capable of supporting a growing ecosystem of manufacturers, dealers, service partners, and drivers, while maintaining high levels of availability.

The company has also connected its manufacturing plants through resilient networks designed to minimise operational disruption. Improvements in system reliability and network stability have translated into lower operating costs and reduced downtime.

What comes next

Looking ahead, Tata Motors plans to deepen its use of AI across manufacturing and service operations. Areas under consideration include intelligent devices on the factory floor, digital representations of dealer operations to improve service efficiency, and tighter integration of data across the enterprise. As data protection requirements evolve, automotive companies will need to balance innovation with stronger governance and security.

Pandey suggests that Tata Motors views this transition as structural rather than cyclical. The company is not simply adopting new tools. It is adjusting to a world in which software, data, and intelligence increasingly determine how vehicles are built, sold, and used.

For an industry long defined by metal and machines, that represents a transformational change. ■

Why enterprises must move from experiments to outcomes



Mahendra Upadhyay, CIO, Broadcast Audience Research Council (BARC), on scaling AI responsibly and driving enterprise transformation

By **Musharrat Shahin** | editor@cioandleader.com

In conversation with CIO&Leader, Mahendra Upadhyay, CIO at BARC India, shares how enterprises are moving beyond AI experimentation toward real business value, emphasising responsible AI, enterprise resilience, and

customer-first digital strategies. He also highlights the growing role of AI agents, talent transformation, and trust-driven ecosystems in shaping the next phase of enterprise innovation.

For much of the past two years, AI has dominated boardroom conversations, product announcements, and investment strategies. Yet beneath the noise of new models and tools, a quieter shift is taking place inside enterprises. Technology leaders are beginning to move from experimentation to real-world value. The focus is no longer on proving that AI works but on proving that it can solve meaningful business problems at scale.

As Mahendra Upadhyay explains, the last year was largely about discovery. Organisations rushed to announce AI initiatives, test proofs of concept, and explore possibilities. But the next phase will demand something far more difficult: measurable outcomes. “The time of announcements and MVPs is over,” he says. “Now the real question is whether AI can solve big business problems.”

From announcements to realisation

The early surge of generative AI was driven by rapid innovation from technology leaders and the explosive popularity of large language models. Companies quickly announced collaborations, integrated AI assistants into productivity platforms, and experimented with automation in workflows.

But early excitement also created confusion. Organizations rushed to adopt AI without fully understanding where it would create sustainable value.

According to Upadhyay, the coming year will shift the focus toward practical implementation. CIOs will remain responsible for three enduring priorities: delivering business value, reducing technology debt, and ensuring visibility into how technology supports competitive advantage. But a fourth priority is now emerging.



Mahendra Upadhyay
CIO, BARC India

Alongside cost reduction and value creation, enterprises must now think about how responsible AI becomes part of their work culture,” he explains. Protecting intellectual property, securing data, and ensuring ethical deployment are no longer optional considerations. They are becoming core governance requirements.

The rise of agentic AI

Earlier waves of automation were dominated by robotic process automation and workflow orchestration. These systems followed predefined instructions to complete repetitive tasks. AI agents introduce a different paradigm. Instead of executing rigid scripts, they can analyze data, interpret context, and respond dynamically. This shift is often described as the move from automation to autonomy.

Customer-facing chatbots, for instance, are evolving into intelligent assistants capable of understanding customer intent, personalizing responses, and recommending products or services. By combining insights from CRM systems, sentiment analysis, and behavioral data, enterprises can create highly personalized interac-

tions that drive both customer satisfaction and revenue growth.

For many organizations, this represents a new opportunity to rethink how technology engages customers. Rather than stitching together disconnected digital journeys, companies can design platforms that deliver seamless, experience-driven interactions.

Unlocking value through AI investments

As enterprises move beyond experimentation, the focus of AI investment is also changing. Instead of isolated proofs of concept, organizations are looking for measurable returns. AI initiatives are increasingly evaluated based on their ability to reduce operational expenditure, accelerate decision-making, and create new revenue streams.

For instance, enterprises are beginning to examine where AI can optimize large cost centers. If technology investments can reduce operating expenses over two or three years while improving efficiency, the business case becomes far more compelling. The real challenge lies in scaling these benefits.

Many organizations have deployed AI in small pilots but struggle to translate those pilots into enterprise-wide solutions. Achieving scale requires not only better algorithms but also stronger collaboration between business and technology teams.

Enterprise resilience in the age of AI

The expansion of AI also introduces new risks, particularly as organizations rely on data-intensive systems and distributed infrastructure.

Historically, disaster recovery strategies focused on metrics such as Recovery Point Objective (RPO) and Recovery Time Objective (RTO). But AI-driven systems require a broader perspective. Enterprises

must now think about resilience at the organisational level rather than the application level.

“When AI operates across massive datasets and interconnected systems, resilience cannot be limited to individual applications. It has to become enterprise-wide.” Upadhyay noted. This shift requires new governance frameworks, faster recovery mechanisms, and stronger accountability structures.

Cybersecurity also becomes more complex as AI systems expand the digital attack surface. Instead of simply reacting to incidents, organizations are increasingly using AI-driven detection to identify threats and respond proactively.

The goal is no longer just mitigation. It is anticipation.

The strategic shift: Customer first, AI first

For years, digital transformation initiatives revolved around cloud-first and mobile-first strategies. While those priorities remain important, they are now being complemented by a new philosophy.

Customer-first. AI-first: This approach recognizes that technology must adapt to customer behavior rather than forcing customers to adapt to technology.

Whether customers interact through mobile apps, websites, or digital platforms, enterprises must deliver consistent, personalized experiences across every touchpoint. AI enables organizations to analyze context in real time and respond with tailored interactions that feel seamless to the user.

In practice, this means data must flow freely across systems while maintaining strong security and privacy protections. The result is a more synchronized digital ecosystem, where customer interactions become more intuitive and responsive.



Only when AI initiatives are aligned with core business objectives can they move from experimentation to transformation.

Talent: The real competitive advantage

Despite the technological excitement surrounding AI, Upadhyay believes the most underestimated challenge is human capability.

Enterprises cannot simply hire AI expertise overnight. Instead, they must invest in learning, reskilling, and cultivating talent within their own organizations.

“We will not find the right talent with the click of a button. We have to build it,” he emphasised.

This requires a cultural shift. Employees must be encouraged to experiment, learn new tools, and collaborate with AI rather than fearing that automation will replace their roles. In many cases, AI will eliminate repetitive tasks but create new opportunities for architects, designers, and strategic thinkers.

Organizations that frame AI as a partner rather than a threat are

far more likely to build a workforce capable of thriving in the new digital landscape.

Building trust

Ultimately, the long-term success of AI adoption depends on trust.

Enterprises must create ecosystems where customers, employees, and partners feel confident that technology is being used responsibly. That means balancing innovation with safeguards such as data privacy, secure infrastructure, and transparent governance frameworks.

“If we can build a trustworthy ecosystem around privacy, personalization, and secure computing. The opportunities ahead will be enormous.” Upadhyay said.

Artificial intelligence may be redefining how enterprises operate, but its success will depend on a simple principle. Technology must empower people. ■



Viraj Deshpande
Sr VP & CIO, Petrochemicals,
Reliance Industries

How CIO accountability is being rewritten for 2026

Viraj Deshpande, Sr. VP & CIO of Petrochemicals at Reliance Industries, shares how expectations, metrics, governance, and relationships are being redefined for 2026 and beyond

By **Musharrat Shahin** | editor@cioandleader.com

Reliance’s Intelligence Manifesto, articulated by our Hon’ble Chairman, sets a clear ambition, to become India’s first AI-native enterprise—one that systematically converts data into insights, insights into actions, and actions into customer delight.

Understanding every customer with unmatched depth and empathy, creating experience so intelligent and effortless, they feel like magic. Intelligence is our power; India is our purpose.

Over the last three years, the CIO mandate has undergone a fundamental shift. Technology is no longer evaluated as an enabler, a cost center, or even a transformation program.

It is now assessed as a direct driver of enterprise outcomes, with explicit accountability for value

realization, risk, resilience, and speed of execution.

In asset-intensive sectors such as petrochemicals, this shift is even more pronounced. It is no longer asked what systems were delivered; they ask what business outcomes moved, what risks were reduced, and what decisions became faster and more reliable as a result of technology investments.

Drawing on the AI-Native Petrochemicals execution framework and outcome model, that reflects a real-world CIO perspective on how expectations, metrics, governance, and relationships are being redefined for 2026 and beyond.

CIO&Leader: In 2026, what are the top business outcomes businesses expect from technology, and how have those expectations

shifted over the past three years?

VIRAJ DESHPANDE: In 2026, expectation from technology is to deliver measurable business outcomes, not abstract digital maturity. The shift over the last three years is clear:

- **From efficiency to value creation:** Technology is now expected to improve EBITDA, ROCE, working capital, and margin resilience, not just reduce cost.
- **From automation to decision superiority:** Expectation is for better, faster, and more consistent decisions across pricing, asset performance, supply chain, and capital allocation.
- **From digitization to AI-native execution:** The expectation is no longer digitized workflows, but closed-loop, AI-executed workflows with human supervision.

In petrochemicals, this translates into outcomes such as:

- Higher on-stream factors and lower unplanned downtime
- Faster grade introduction and portfolio optimization
- Improved OTIF with lower inventory
- Disciplined pricing, credit, and working-capital control
- Improved customer value, operational excellence, and shareholder returns

From digitized processes to AI-executed workflows

The expectation is no longer dashboards or reports but closed-loop workflows where AI executes repeatable operational and commercial decisions within defined safety, DoA, and regulatory guardrails—under human supervision.

Technology is now judged by its ability to move these needles consistently, across cycles..

CIO&Leader: How are CIOs demonstrating measurable value from digital, AI, and transformation investments to the board, and which metrics truly matter today?

VIRAJ DESHPANDE: CIOs no longer defend spend using system uptime or adoption metrics. Value is demonstrated through workflow-linked KPIs, plant-level and P&L-visible outcomes

What resonates with executive leaderships today:

- **Outcome-linked metrics**, such as margin per ton, inventory days, yield improvement, or reduction in manual intervention hours
- **Decision cycle-time reduction**, especially in planning, logistics, pricing, and turnaround execution
- **Risk-adjusted value**, showing not just upside but volatility reduction and downside protection

The expectation is no longer dashboards or reports but closed-loop workflows where AI executes repeatable operational and commercial decisions within defined safety, DoA, and regulatory guardrails—under human supervision.

In an AI-native model, each canonical workflow (e.g., Order-to-Delivery, Asset Performance, Billing & Credit) has

- A clear business owner
- Codified decision rights (DoA-as-code)
- Embedded policy guardrails
- Quantified impact on EBITDA, ROCE, safety, or working capital

This allows CIOs to report value at the level, where outcomes are appreciated, not platforms..

CIO&Leader: Where do expectations align with reality, and where do execution complexity and trade-offs create the biggest challenges on the ground?

VIRAJ DESHPANDE: There is strong alignment between Business Leadership and CIOs on what must be achieved. The challenge arises in how fast and how cleanly it can be executed.

Key on-ground challenges include:

- **Legacy fragmentation** across plants, functions, and systems that resists end-to-end workflow ownership
- **Change fatigue** when transformation is layered on top

of existing operating models instead of rewiring them

- **Data readiness gaps**, where AI ambition outpaces data quality and governance maturity
- **Risk aversion around autonomy**, especially in safety-critical environments

This is why parallel-run (“make before break”) models are essential. AI-led execution must prove superiority on safety, reliability, stability, and outcomes before authority is shifted not as a leap of faith, but as an evidence-based transition..

CIO&Leader: Which technology risks are now viewed as enterprise risks, and where are accountability boundaries still unclear between the CIO, CISO, and business leaders?

VIRAJ DESHPANDE: Several technology risks are now firmly viewed as enterprise risks:

- AI decision integrity and explainability
- Data quality and lineage
- Cloud cost volatility
- OT cybersecurity and plant safety interlocks

However, accountability boundaries remain blurred. CIOs own platforms and data, CISOs own cyber defense, but business leaders own outcomes. The most effective organizations make this explicit: technology risk governance is shared, with clear escalation paths and codified limits..

CIO&Leader: How has the cybersecurity conversation evolved from protection to enterprise-wide resilience, recovery, and accountability?

VIRAJ DESHPANDE: Cybersecurity has evolved from a perimeter defense conversation to an enterprise resilience mandate.

The expectation is:

- **Business-continuity assurance**, not just breach prevention



CIOs no longer defend spend using system uptime or adoption metrics. Value is demonstrated through plant-level and P&L-visible outcomes.

- **Time-to-recover metrics**, especially for plants, logistics, and order fulfillment
 - **Clear accountability** for cyber-physical risks in OT, IT, and AI systems
- In AI-native enterprises, security is embedded as policy-as-code, not manual enforcement. Autonomous workflows cannot bypass safety, regulatory, or financial guardrails. As autonomy increases, governance must become stronger not looser.

CIO&Leader: How are boards reassessing cloud costs, ROI, and modernization priorities, and what tensions does this create for long-term digital strategy?

VIRAJ DESHPANDE: Cloud econom-

ics are being assessed with sharper scrutiny. The question has shifted from cloud first to value first.

- Imperatives considered are:
- Short-term cost control conflicts with long-term modernization
 - Lift-and-shift models fail to deliver promised agility
 - AI workloads demand scalable platforms but lack clear outcome ownership

Successful CIOs anchor cloud investments to specific workflow economics — for example, how faster planning cycles or predictive maintenance reduce working capital or downtime. Cloud is no longer a strategy; it is an execution enabler.

CIO&Leader: What do CIOs need more from boards to deliver outcomes successfully, clearer direction, investment discipline, patience, or shared risk ownership?

VIRAJ DESHPANDE: CIOs do not need more enthusiasm for technology. They need:

- Clear outcome priorities, not shifting digital agendas
- Investment discipline, tied to measurable value realization
- Patience for foundational rewireing, especially data, governance, and operating-model change
- Shared ownership of risk, particularly in AI-led decisioning and autonomy

When business treat transformation as an enterprise shift — not an IT program — execution accelerates dramatically.

CIO&Leader: What is the most underestimated technology risk or opportunity for 2026–27, and how do expectations differ across sectors?

VIRAJ DESHPANDE: The most underestimated opportunity is AI-native execution at scale — not pilots, copilots, or dashboards, but AI systems that run the business within defined boundaries.

Conversely, the most underestimated risk is partial transformation: adopting AI without rewriting workflows, governance, and accountability. This creates complexity without compounding value.

Sector expectations vary. Asset-heavy industries focus on reliability and safety; consumer sectors focus on personalization and speed. But across sectors, the winners will be those who treat AI not as a tool — but as a new operating system for the enterprise. ■



Ashish Thakur
CIO, Cummins India

CIOs are now accountable for outcomes, not infrastructure

Ashish Thakur, CIO, Cummins India on what boards really expect from technology in 2026

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

As enterprises move from digital transformation to decision transformation, CIOs are under growing pressure to prove measurable business outcomes from technology investments. In this conversation with CIO&Leader, Ashish Thakur unpacks how board expectations have shifted from cloud adoption and pilots to revenue impact, resilience, and decision quality in an AI-first enterprise. He explains why use-case-driven AI, integrated digital cores, and outcome-linked metrics now define CIO credibility and why boards must move from passive oversight to shared risk ownership to unlock real transformation.

CIO&Leader: In 2026, what are the top business outcomes businesses expect from technology, and how have those expectations shifted over the past three years?

ASHISH THAKUR: Over the past three years, technological expectations have evolved significantly. Businesses have shifted from founda-

tional transformation like cloud migration and remote work enablement, to leveraging technology as a revenue driver. Autonomous systems, AI-assisted processes, and real-time predictive capabilities such as scenario planning, demand sensing, and supplier risk modelling have become essential. The focus is now on pre-empting disruptions rather than reacting to them. Customers demand faster innovation cycles, personalized products, and agile market response. Integration across PLM (Product Lifecycle Management), ERP (Enterprise Resource Planning), and MES (Manufacturing Execution System) is imperative, enabling seamless workflows from design, simulation, commissioning and service. CIOs must prioritize use-case-driven AI adoption, ensuring technology delivers measurable value, beyond just tools and systems.

CIO&Leader: How are CIOs demonstrating measurable value from digital, AI, and transformation

investments to the board, and which metrics truly matter today?

ASHISH THAKUR: CIOs today are increasingly expected to demonstrate measurable value from digital, AI, and transformation investments, shifting the conversation from mere technology adoption to clear business impact. Boards and investors now prioritize ROI that aligns with strategic goals and tangible outcomes. To address this, CIOs are linking technology investments directly to core business KPIs, enabling a clear correlation between tech initiatives and business performance. They are also leveraging frameworks such as OKRs (Objectives & Key Results) and Value Stream Mapping to track and showcase progress against strategic objectives. Benchmarking against industry standards further builds credibility by contextualizing performance.

Additionally, storytelling with data has become a critical skill for CIOs, as they must present their narratives in a way that highlights business value achieved, rather than

focusing solely on the technology itself. This approach ensures technological initiatives are not viewed as abstract concepts but as drivers of measurable business success.

CIO&Leader: Where do expectations align with reality, and where do execution complexity and trade-offs create the biggest challenges on the ground?

ASHISH THAKUR: Reality sets in when technology moves from concept to deployment, since success relies not just on implementation, but on adoption and effective change management. Success hinges on shared responsibility between business leaders and CIOs with clear ownership and pre-defined success criteria guiding initiatives. These criteria provide a framework for consistent execution and progress tracking. A significant challenge lies in understanding business process complexity. Many initiatives & projects fail because teams underestimate real-world scenarios and exceptions. It's often these exceptions, rather than standard processes that disrupt execution. Addressing these complexities with thorough planning and adaptability is essential to bridge the gap between expectations and on-the-ground realities.

CIO&Leader: What do CIOs need more from boards to deliver outcomes successfully, clearer direction, investment discipline, patience, or shared risk ownership?

ASHISH THAKUR: In 2026, the role of CIOs has evolved from asking for permission to transform, to delivering measurable outcomes like growth, resilience, productivity, and trust, despite rising expectations and reduced tolerance for errors. The success of these efforts depends as much on the board's behavior, governance, and spon-

“Leaders must shift focus from counting AI initiatives to ensuring AI enhances decision quality. The priority now is to use AI to drive smarter, faster, and more responsible decisions, while managing risks such as bias, opacity, and overreliance.”

sorship along a similar scale on technology capabilities. CIOs need boards to provide clearer direction, including explicit articulation of business outcomes that matter most, and disciplined investment strategies focused on high-impact initiatives rather than spreading resources thin. Additionally, patience is essential, ensuring funding supports end-to-end value realization rather than stopping at pilot programs.

Above all, shared risk ownership by the board, through active sponsorship of cross-functional initiatives, clear prioritization of speed versus risk or growth versus cost, and support in dismantling organizational silos, forms the foundation for achieving transformative outcomes. This collaborative approach fosters alignment and ensures CIOs are equipped to drive meaningful, measurable business results.

CIO&Leader: What is the most underestimated technology risk or opportunity for 2026–27, and how do expectations differ across sectors?

ASHISH THAKUR: Technology risks like AI bias, data integrity, and decision-making distortions are now enterprise risks, with accountability boundaries between CIOs, CISOs, and business leaders often unclear. While AI's promise lies in automation and cost reduction, its deeper impact shaping and accelerating decisions, poses significant risks if unchecked. For example, AI systems expected to improve efficiency may also unlock hidden supply chain costs or margins through quick decisions, yet these outcomes depend on proper oversight. Leaders must shift focus from counting AI initiatives to ensuring AI enhances decision quality, asking: How is AI influencing critical business decisions, and are the results measurably distinct as a comparative analysis of decision-making quality levels?

CIO&Leader: What is the most underestimated technology risk or opportunity for 2026–27, and how do expectations differ across sectors?

ASHISH THAKUR: AI's transformative influence on organizational decision-making appears as among the most underrated technology risks towards the future trajectory. AI's real impact isn't in automation or cost cutting alone, it lies in how it shapes, accelerates, and improves critical decisions. Beyond efficiency gains, AI can reveal hidden opportunities such as supply chain cost reduction, margin optimization, and smarter planning. Success is not about the number of AI projects, but about better decision quality. Leaders should ask: How many critical decisions are AI-enabled, and are outcomes measurably better? The priority now is to use AI to drive smarter, faster, and more responsible decisions, while managing risks such as bias, opacity, and overreliance. ■



Tirthankar Lahiri

SVP, Mission Critical Data and AI Engines, Oracle

You can't bet your business on AI you can't verify

In this exclusive interaction, **Tirthankar Lahiri, SVP, Mission-Critical Data and AI Engines, Oracle**, explains why CIOs must build verifiability—not just intelligence—into critical systems

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

As enterprises move AI from pilots to production, a harder question is emerging inside boardrooms and CIO offices: how do you add intelligence to core systems without breaking trust, governance, or uptime?

The problem is no longer whether AI works. The question is whether AI can be trusted to operate in mission-critical systems that run banks, telecom networks, insurers, and national infrastructure. The shift underway is architectural, not experimental.

In this conversation with Jatinder Singh of CIO&Leader, Tirthankar Lahiri, Senior Vice President, Mission Critical Data and AI Engines at Oracle, argues that the future of enterprise AI will not be defined by speed or scale alone, but by verifiability, structural constraints, and security anchored at the data layer. From why AI must live inside the database to why “guardrails” are not enough for agentic systems,

Lahiri lays out a pragmatic blueprint for CIOs trying to move from automation to accountable delegation.

CIO&Leader: As AI shifts from experimentation to mission-critical workloads, what global changes do you see by 2026 in how enterprises are architecting data platforms that must balance AI innovation with determinism, reliability, and regulatory compliance?

TIRTHANKAR LAHIRI: By 2026, the biggest shift will be architectural. Enterprises are realizing that AI cannot sit outside the system of record. Moving data to separate AI engines introduces governance gaps and security fragmentation.

We are seeing three global patterns:

- **Security anchored at the data layer:** AI must obey the same access controls as human users.
- **Verifiability built into workflows:** Outputs must be audit-able and reproducible.

- **Bounded intelligence:** AI must operate within defined execution paths, not open-ended reasoning loops.

Innovation is no longer about experimentation. It is about embedding intelligence inside deterministic infrastructure.

CIO&Leader: Why do you believe AI must be engineered directly into the data platform, and what risks arise when AI is layered outside core enterprise systems?

TIRTHANKAR LAHIRI: What we are doing at Oracle is focused within the database organization. The strategy is simple: you should not have to move your data elsewhere to run AI. Many older vector database products required moving data to a separate system for AI workloads. That creates security risks because once data leaves the corporate database, you lose governance and control.

For us, everything must remain in the same place. Enterprises can only deploy solutions they can not only trust but also verify. Trust implies faith. In enterprise systems, faith is not enough. You must be able to verify the output.

It starts with security embedded deep in the data layer. We call this Deep Data Security. When AI accesses data, it must follow the same rules as a human user. AI cannot bypass security controls. If security lives only in the application layer, AI can bypass it and access data directly. Security has to exist at the source.

We also focus on verifiable design. We call it Generative Creator Development. Generative AI is not one technique. It is a collection of approaches to ensure outcomes are trustworthy and editable.

If I receive an output, I should be able to review it and decide what I want and what I do not want. AI-generated code and SQL are risky because you may not know what they will ultimately do.

So we rely on trusted data APIs that prevent arbitrary access to tables and columns. They expose complete business objects and ensure any change preserves business validity. The same applies to agentic AI. Every interaction must be verified. Guardrails alone are not sufficient.

CIO&Leader: You already touched on GenAI and described Database 26ai as AI native. How does that fundamentally change capabilities such as vector search and natural language interaction when these are built directly into mission-critical applications?

TIRTHANKAR LAHIRI: Vector search brings AI into any application by expanding how data can be queried. I can ask, “Find me the top-selling products.” That runs on any version of Oracle Database. But now I can



“AI should act like a clerk, not a judge. The final decision must remain human.”

also ask, “Find me the top-selling products that are similar to this product.” Similarity is no longer a traditional database construct. It is semantic. With an AI-powered database, that becomes possible. Traditionally, databases excel at value-based operations: precise filtering, aggregation, and computation. With AI and vector search, you can search based on semantic similarity, which traditional databases could not do. That is foundational to modern AI systems, and that is why it is transformational.

Vector search itself may not be unique to Oracle, but Oracle’s strength lies in combining it with the database’s other capabilities. You can take complex applications in banking, financial services, or telecom and embed AI simply by adding vectors as another search dimension. That integration is what makes the shift fundamental. It also enables natural language interaction.

Looking ahead, I believe most systems will primarily use natural language interfaces. SQL and pro-

gramming languages may become like assembly language, low-level outputs of higher-level conversational interactions.

But human language is ambiguous. If someone asks, “Find me the top-selling products,” the system should clarify: “Do you mean by revenue or by number of units?” The shift from an ambiguous question to a precise, deterministic query is critical.

This will be multi-stage. AI interprets the initial question, refines it into a more specific question, and presents it back for confirmation. For example: “Do you want the products with the highest units sold in Maharashtra last week?” Once the user confirms, the translation to SQL becomes straightforward and deterministic.

So AI is the bridge between natural language and structured business data through semantic matching. That is where Oracle’s strength lies, given the scale and importance of the business data repositories it manages.

CIO&Leader: With AI-specific regulations emerging alongside DPDP, how should enterprises design AI systems to ensure auditability, accountability, and verifiability rather than blind trust?

TIRTHANKAR LAHIRI: For me, it has to be verifiable. AI cannot operate like a black box. Every action must be auditable and traceable. There must be reasoning behind why an answer was selected, and a human must always be able to edit it.

I prefer a multi-stage generation process. Instead of producing one long, unreviewable output, AI should generate structured responses that can be checked. Even if it runs automatically, everything must be logged so someone can go back, review what happened, and correct it.

The confusion comes from consumer AI. People think they can use tools like ChatGPT to summarize business data and rely on it. You should not bet your business on that. AI is nondeterministic. It may be right most of the time, but that is not enough for enterprise systems.

AI is like a smart teenager. Often correct, always confident, and ready with an answer, even when it is wrong. That can be useful, but it requires supervision. Enterprise AI has to operate within workflows that support editability, accountability, and retry. I think regulation will evolve around that reality. Verification will always be required.

CIO&Leader: Some of Oracle's large customers, especially in banking and other regulated sectors, still value a predictable database over an intelligent one. How do you address concerns that introducing AI into the database could create opacity, governance challenges, or new security risks?

TIRTHANKAR LAHIRI: First, security must live in the data layer. That is

“Innovation is no longer about experimentation. It is about embedding intelligence inside deterministic infrastructure.”

non-negotiable. You cannot secure AI only at the application layer. Data is the foundation, so secure the foundation.

Second, introduce AI carefully. Start with internal use cases before exposing them to customers. Many banks are experimenting with internal chatbots to improve employee productivity. That is a lower-risk entry point.

You should not use AI for adjudicative workflows. An AI should not approve mortgages or claims. But it can assist a human agent by drafting questions, recommending additional information, and helping structure an application. The final decision must remain human.

AI should act like a clerk, not a judge. It can organize and assist, but it should not independently approve or reject critical decisions. That is how you add intelligence without compromising governance.

CIO&Leader: As intelligence spreads across the stack, does it compound value or compound risk? What controls are essential at the engine level?

TIRTHANKAR LAHIRI: It can do both.

Intelligence compounds value when it improves efficiency inside bounded workflows. It compounds risk when it operates without structural limits.

Essential controls include:

- Hard access restrictions, not soft guardrails
- Deterministic execution pathways
- Full logging of AI interactions
- Human override capabilities

AI should operate like a train on tracks. It may choose among predefined routes, but it cannot leave the rails.

CIO&Leader: From an engineering level standpoint, what guardrails do you think are critical to prevent unintended or emergent behavior in production systems?

TIRTHANKAR LAHIRI: I personally do not believe in guardrails. They give a false sense of security. Guardrails are still suggestions. Think of it like telling my teenage child, “Please don’t go into that room.” That is my guardrail. I can be pretty sure someone will go into the room. So what do you do? You lock the room.

You cannot expect AI to follow rules just because you told it to. You have to engineer trust into the core architecture. Guardrails are, at best, documentation of what the rules should be. What you really need are hard rails. Hard tracks you cannot deviate from. You can choose a direction, you can go to A or B, but you cannot go in an arbitrary direction.

That is the only way to build a secure, verifiable enterprise AI.

CIO&Leader: With the rise of vector databases and specialized AI data stores, do you think the traditional RDBMS model is moving away?

TIRTHANKAR LAHIRI: No, I think the opposite. My prediction is that vector databases were a flash in the pan because of generative AI. Companies quickly discovered that once you move data out of a relational database, you lose security and governance.

In the next five years, every major database will support vectors natively. We will not even discuss vector indexes in five years because they will be like the indexes everybody has.

When that happens, standalone vector databases will struggle. They are not bad products, but they are limited in scope. They can perform semantic similarity searches, but they do not support sophisticated relational business queries or filters. Once you export data to them, your security controls are no longer in place. You have to reinvent them, and those databases lack decades of enterprise security maturity.

Oracle, for example, has deep security capabilities, like only letting me see employees in my organization from a table. A vector database lacks that sophistication. If I run a similarity search there, it may run across everyone.

I think the benefits of vector databases will move into the traditional relational databases and expand what they do. But I do not think vector databases will have a long-lived run as document databases did. I would be surprised if, five years from now, many of them still exist.

CIO&Leader: Looking at India specifically, where do you see maximum demand for AI cloud solutions coming from? What are you seeing among Oracle's largest mission-critical customers and their key pain points?

TIRTHANKAR LAHIRI: My largest mission-critical customers are focused on core systems like core banking and core telecom, the systems their businesses depend on.

A lot of them are enabling RAG and interactive workflows on those systems to answer questions like, "Which customers are similar to this customer?" or "What service request is similar?" That is a very

"Enterprises can only deploy AI they can not just trust, but verify. Faith is not enough in mission-critical systems."

common use case. Before taking a new service request or ticket, you first check whether any existing tickets match the symptoms. You can do that on the production service request database.

Service request databases are huge. For a large telecom, you could have millions filed daily. Earlier, you would just accept the request and process it. Now you can do a first diagnosis: "This looks similar to that issue. It may be the same problem."

That creates massive operational efficiency. Instead of routing it to another rep and duplicating work, you can flag likely duplicates and speed up resolution. These core systems will be augmented with AI to drive better operational efficiency.

CIO&Leader: Is adoption of AI-enabled databases accelerating with trust, or are they deploying cautiously due to risks?

TIRTHANKAR LAHIRI: People are cautious. Not just in India. They are still exploring what they can do without harming anyone. That is the right way to start. Find use cases with no risky side effects that improve efficiency and give measurable ROI, but without taking over workflows.

Caution is the name of the game in enterprise, especially mission-critical systems. I do not see agentic workloads proliferating immediately on core systems. I see adjunct workflows that improve efficiency. Over time, agentic workflows will

become more common as we get comfortable with them. In banking, especially given the legislation and regulations to follow, cautious adoption is the right approach.

CIO&Leader: What is the single biggest driver, cost, speed of innovation, efficiency, or something else?

TIRTHANKAR LAHIRI: Efficiency. Everybody wants to reduce costs and improve operations. That is a big driver.

As in the service request example, one customer saw a 60-70% reduction in service requests filed because they could detect duplicates quickly. That reduced the workload for the team managing them. These are the early use cases customers want: reduce costs and improve turnaround time

CIO&Leader: What leadership or architectural trade-offs do CIOs face introducing intelligence into mission-critical databases without compromising trust, uptime, or accountability?

TIRTHANKAR LAHIRI: Ideally, you should not make trade-offs. Trade-offs imply compromise, and most CIOs do not want to compromise on system integrity or verifiability. They want to add AI seamlessly without introducing risk or vulnerabilities.

The ideal outcome is adding AI without compromising integrity.

CIO&Leader: How critical is upskilling to manage risks like data leaks from employees using rogue AI tools like ChatGPT?

TIRTHANKAR LAHIRI: Upskilling is essential. AI won't replace humans, but humans who master AI will replace those who don't. You can't stop large workforces from using ChatGPT independently; education on safe AI practices is absolutely key. ■



Hemant Tiwari

Managing Director & VP – India and SAARC, Hitachi Vantara

Why CIOs must rethink their data stack in 2026?

Hemant Tiwari, Managing Director & VP – India and SAARC, Hitachi Vantara highlights the data strategy within defined jurisdictions

By **Musharrat Shahin** | editor@cioandleader.com

In conversation with CIO&Leader, Hemant Tiwari, Managing Director & VP – India & SAARC, Hitachi Vantara, shares how Indian CIOs can build sovereign, hybrid data foundations aligned to DPDP, RBI and MeitY mandates. He unpacks how to modernise legacy storage, enable AI-ready infrastructure, and embed zero trust at the data layer. The discussion also explores cyber resilience, edge-to-cloud architectures, and governance at scale.

CIO&Leader: How is Hitachi Vantara tailoring its global data platform strategy to align with India's unique priorities such as data sovereignty under the DPDP Act, "Make in India", and sectoral mandates from RBI and MeitY?

HEMANT TIWARI: India's regulatory and digital priorities demand a data strategy that is both globally robust and locally precise. At Hitachi Vantara, we design platforms that enable data sovereignty by keeping sensitive data governed,

protected, and accessible within defined jurisdictions. Our core data infrastructure and hybrid cloud strategy deliver policy-centric controls, robust governance, and security posture that support compliance demands across industries. At the same time, we invest in local capabilities and partner engagement to support digital growth initiatives and innovation across government and enterprise sectors. This approach helps CIOs build a future-ready foundation aligned to national and regulatory priorities while enabling scalable outcomes.

CIO&Leader: Can you share some examples where Hitachi Vantara has co-developed or localized data solutions with customers to address India-specific challenges like vernacular data, rural edge connectivity, or regulatory complexity?

HEMANT TIWARI: Our customer engagements in India focus on practical outcomes that reflect local

needs. We work with partners and customers to tailor hybrid data solutions that cope with distributed operations, vernacular and large-scale unstructured data, and environments with variable connectivity. We also build localized proof-of-concepts and workshops with partner ecosystems, ensuring that regulatory complexity is translated into consistent compliance and governance frameworks within the data layer. This collaborative approach allows CIOs to deploy resilient data platforms that remain relevant to India's diverse operational landscape while leveraging global best practices.

For example, with Malayala Manorama, one of India's largest regional media groups, we worked closely with their technology teams to modernize data infrastructure for managing high volumes of vernacular and unstructured content across print, digital, and mobile platforms. The solution was tailored to support rapid content ingestion,

scalable storage, and secure access across distributed newsrooms, enabling faster publishing cycles while maintaining strong data governance and long-term archival integrity.

In the BFSI and IT services sectors, our engagements with an Asia-based bank operating in India and Infosys focused on building resilient, compliant hybrid data platforms that align with local data residency, risk, and audit requirements. We worked closely with customers to design architectures that ensure high availability, secure recovery, and consistent data governance across distributed delivery environments, helping organizations confidently scale digital services while adapting global best practices to India's regulatory and operational landscape.

CIO&Leader: Many organizations still rely on fragmented, aging storage systems. What's your recommended path to modernize this infrastructure while minimizing business disruption?

HEMANT TIWARI: The journey to modernization begins with consolidation, visibility, and prioritizing non-disruptive transformation. We recommend leaders start by unifying fragmented storage and data estates into intelligent data platforms that simplify operations and automate management. Early adopters benefit from predictable performance, reduced complexity, and improved resilience without disrupting business-critical workloads. By taking a phased approach, leveraging automation and policy-based governance, organizations can modernize steadily while unlocking the performance, scalability, and flexibility needed for digital-era workloads.

CIO&Leader: With AI becoming boardroom imperative, how does



Hitachi Vantara ensure its storage and data platforms deliver the performance, reliability and integration needed to power scalable, responsible AI workloads?

HEMANT TIWARI: AI success starts with a reliable, high-performance data foundation. With platforms designed to unify block, file, and object data and native integrations across hybrid environments, organizations gain consistent performance and governance. Our hybrid cloud and data management capabilities help CIOs move from experimentation to production with enterprise-grade reliability, enabling responsible AI adoption that is governed, secure, and optimized for real business value.

CIO&Leader: Given the sharp rise in ransomware targeting Indian enterprises, how does your data stack provide protection and resilience?

HEMANT TIWARI: We see ransomware today not as a perimeter failure, but as a data survivability challenge. The way our data stack is designed, we aim to predict breach

and focus on guaranteed recoverability, immutability, and business continuity. At the core of our approach is a zero-loss recovery architecture that combines immutable snapshots, air-gapped backup copies, and continuous data integrity validation. In case the primary systems are impacted, we focus on supporting enterprises in restoring clean data copies in a short span of time, helping to minimize operational disruption.

What further distinguishes our approach is that resilience is supported through our Cyber Resilience Guarantee, which outlines recovery objectives when recommended best practices are followed. For Indian enterprises navigating increasing regulatory scrutiny and the financial impact of downtime, this helps position cyber resilience as part of a broader business risk management strategy rather than solely an IT consideration.

CIO&Leader: Beyond perimeter security, how are zero-trust principles embedded into your data infrastructure?

HEMANT TIWARI: At Hitachi Vantara, we believe zero trust must extend all the way to the data layer, not stop at identity or network controls. Our approach assumes that internal networks can be just as hostile as external ones, which is why we embed Zero Trust Architecture directly into the data infrastructure, ensuring that security travels with the data rather than remaining perimeter bound. Every access request, whether from a user, application, cloud workload, or edge device, is treated as untrusted by default and continuously verified based on identity, role, and contextual signals before any interaction with data is allowed.

From an implementation standpoint, we enforce zero trust through fine-grained access controls, least-privilege permissions, and strict separation between production, backup, and recovery environments, all with full auditability, even privileged users. We apply micro-segmentation to prevent lateral movement and contain breaches to the smallest possible footprint, while data-centric encryption protects information both at rest and in transit across hybrid and multi-cloud environments. By embedding these controls directly into our storage and data protection layers, we significantly reduce the blast radius of insider threats, credential compromise, and ransomware attacks, the challenges that are increasingly prevalent across Indian enterprises.

CIO&Leader: The DPDP Act mandates purpose limitation, consent management, and data minimization. How can your platform help in the discovery, classification, and enforcement of dynamic usage policies on personal data across hybrid environments?

HEMANT TIWARI: We at Hitachi Vantara view the DPDP Act as a

“This unified approach empowers organizations to scale at their pace, optimize TCO, and retain control over their data estate in a way that meets compliance and performance expectations.”

catalyst for operationalizing data responsibility, not just compliance. Our platform enables automated discovery and classification of personal and sensitive data across structured and unstructured sources, on-premises, in the cloud, and at the edge. Once data is classified, we help enterprises dynamically enforce usage policies throughout the data lifecycle. This ensures data is used only for defined purposes, retained only as long as required, and governed in line with consent conditions, even as data flows into analytics and AI workloads. Since these controls operate continuously rather than through periodic audits, enterprises gain stronger accountability, reduced regulatory risk, and the confidence to scale data-driven innovation in India’s evolving privacy landscape.

CIO&Leader: Enterprises struggle with fragmented oversight across on-prem, cloud, and SaaS. How does Hitachi Vantara enable a single data governance fabric?

HEMANT TIWARI: Fragmentation can make governance more complex, and we address this through Virtual Storage Platform One (VSP One) and our unified data fabric approach, which aims to provide a true single pane of glass across

the entire data ecosystem. VSP One abstracts the complexity of on-premises, private cloud, and public cloud environments, giving IT leaders a centralized control plane to view, manage, and govern data consistently, regardless of where it physically resides.

Instead of maintaining fragmented policies across AWS, Azure, and local data centers, our data fabric allows enterprises to define governance policies once and enforce them everywhere, ensuring uniform controls across all data assets.

CIO&Leader: Many firms retain data indefinitely due to legal uncertainty. How do your solutions automate data lifecycle management in alignment with Indian sectoral regulations?

HEMANT TIWARI: Across sectors in India, organisations are managing unprecedented data growth while also navigating evolving regulatory expectations. In this context, data often gets retained longer than necessary simply to avoid risk. Our focus is on helping enterprises bring structure and automation to data lifecycle management. Using metadata-driven classification and policy-based controls, data can be governed from the moment it is created. Retention, archival and deletion rules can then be aligned with sector specific needs, whether it is audit requirements in financial services, record management in the public sector, or long term content preservation in media organisations. For example, platforms such as Virtual Storage Platform One and Hitachi Content Platform allow organisations to apply consistent lifecycle policies while maintaining audit readiness and operational clarity.

CIO&Leader: Since enterprise data fuels AI models, how does your governance framework help detect

bias, ensure explainability, and maintain accountability—especially in regulated sectors like BFSI?

HEMANT TIWARI: As AI adoption expands across enterprises, governance becomes a prerequisite for trust. Our approach centres on data quality, lineage and transparency so that organisations have a clear understanding of where data originates and how it is transformed before it is used in AI models. Through solutions such as Hitachi iQ, data can be catalogued, profiled, and traced end to end, which supports explainability and accountability in decision making. In regulated sectors like BFSI, this level of visibility helps teams demonstrate how outcomes are derived and supports internal risk and compliance processes. More broadly, it helps organisations build confidence in AI driven insights across use cases.

CIO&Leader: How portable is data on your platform—and how do you ensure consistent governance regardless of where it resides?

HEMANT TIWARI: Flexibility has become a core requirement for CIOs as organisations adopt hybrid and multi cloud strategies. Our platforms are designed so that governance, security, and lifecycle policies are applied at the data level rather than being tied to a specific location. This means data can move across public clouds, private environments or Indian sovereign clouds while retaining the same controls for access, encryption and compliance. Solutions like Hitachi Content Platform support this portability by enabling a unified view of data across environments, which allows organisations to respond to business or regulatory changes without fragmenting governance practices.

CIO&Leader: With sustainability now a board-level mandate, how do your storage innovations

“AI- and ML-driven anomaly detection continuously monitors data flows and device behaviour, enabling rapid identification of deviations or integrity issues and helping maintain uptime, resilience, and operational reliability.”

help Indian enterprises meet ESG goals and reduce TCO simultaneously?

HEMANT TIWARI: Sustainability and cost management can go together when infrastructure is engineered thoughtfully. Our sustainability reporting outlines how energy efficiency is embedded into modern storage architectures. For example, systems that are ENERGY STAR certified and designed to reduce power consumption through automated efficiency features which help lower carbon footprint while also reducing operational costs for power and cooling. Customers deploying solutions like VSP One Block have experienced meaningful improvements in energy usage and rack space efficiency, which in turn supports broader ESG commitments while positively impacting total cost of ownership. This kind of alignment means technology investments contribute both to financial discipline and sustainability reporting.

CIO&Leader: Beyond uptime and compliance, how should

CIOs quantify the business value of their data infrastructure investments?

HEMANT TIWARI: CIOs should frame infrastructure value in terms that resonate with business goals. Useful KPIs include data availability for analytics and AI workflows, governance maturity scores, time to insight, reductions in redundant or unused data, and operational efficiencies such as lower energy use per terabyte stored. In regulated industries, demonstrating how governance reduces audit cycle times can also be meaningful. For ESG mandates, tracking energy savings or carbon reductions tied to infrastructure platforms creates tangible business value. These metrics help bridge the language between technology performance and organisational outcomes, making it easier to articulate ROI to CFOs and boards.

CIO&Leader: Looking ahead, what three foundational data capabilities should Indian CIOs prioritize today to stay ahead of technological disruption, evolving regulation, and next-gen innovations?

HEMANT TIWARI: Looking at the evolving landscape, three capabilities stand out. First, a unified, metadata-driven governance layer that spans hybrid and multi-cloud environments helps organisations manage risk and compliance consistently. Second, AI-ready data foundations that support transparency, explainability, and bias mitigation enable more reliable and ethical AI adoption. Third, sustainable and resilient infrastructure ensures organisations can grow responsibly while balancing performance, cost, and environmental impact. These foundations help organisations not only respond to current pressures but also support emerging imperatives around AI scale-up and regulatory nuance. ■



Premkumar Balasubramanian
CTO, Hitachi Digital Services

The new CIO scorecard: Revenue, resilience, and AI

Premkumar Balasubramanian, CTO, Hitachi Digital Services, breaks down the metrics and trade-offs boards care about in 2026

By **Musharrat Shahin** | editor@cioandleader.com

In conversation with Premkumar Balasubramanian, CTO of Hitachi Digital Services, we explore the redefined CIO mandate for 2026. As technology moves from a back-office function to a primary revenue driver, Balasubramanian breaks down the "New CIO Scorecard" centered on growth, operational resilience, and the complexities of scaling AI. From closing talent gaps to navigating the "gray zones" of C-suite accountability, this discussion provides a strategic roadmap for turning emerging risks into sustainable enterprise value.

CIO&Leader: In 2026, what are the top business outcomes businesses expect from technology, and how have those expectations shifted over the past three years?

PREMKUMAR BALASUBRAMANIAN: In 2026, businesses expect

technology to deliver the following:

- Revenue growth through AI-driven personalization and automation
- Improve Operational resilience via predictive analytics and autonomous systems
- Rapid innovation with generative AI and low-code platforms
- Sustainability and compliance through green tech and transparent data

Over the last 3 years focus has shifted from:

- **Cost reduction and cloud migration to AI-powered growth and resilience:** Organizations moved from simply cutting IT costs and migrating workloads to cloud toward leveraging AI for new revenue streams, predictive operations, and business continuity.
- **Basic digital transformation**

to intelligent automation and ecosystem integration:

The focus evolved from digitizing processes and moving to digital channels to deploying AI-driven automation and integrating ecosystems for end-to-end efficiency and innovation.

- **Reactive security to proactive, AI-driven risk management:**

Security shifted from responding to breaches after they occur to anticipating threats using AI, implementing zero-trust architectures, and embedding resilience across the enterprise.

CIO&Leader: How are CIOs demonstrating measurable value from digital, AI, and transformation investments to the board, and which metrics truly matter today?

PREMKUMAR BALASUBRAMANIAN: CIOs are tying technology

investments directly to business outcomes, demonstrating how digital and AI initiatives drive revenue growth, improve margins, and enhance customer experience. They emphasize tangible ROI through automation savings, productivity gains, and accelerated innovation cycles, presenting these results in terms the board understands—financial impact and competitive advantage.

The key metrics that matter today are:

- **Revenue impact:** growth attributable to tech initiatives
- **EBITDA uplift:** margin improvement from efficiency
- **Automation rate:** cost-to-serve reduction
- **Innovation velocity:** time-to-market for new products and features
- **Risk posture:** cybersecurity and compliance scores

CIO&Leader: Where do expectations align with reality, and where do execution complexity and trade-offs create the biggest challenges on the ground?

PREMKUMAR BALASUBRAMANIAN: Expectations largely align with reality for delivering automation efficiency, data-driven insights, and improved customer experiences—these outcomes are visible and achievable across industries. On the ground, however, execution is far more complex: teams must integrate modern platforms with legacy systems, close critical talent gaps in AI and cybersecurity, and balance innovation with compliance under tight data-governance standards. Trust in AI remains a hurdle as hallucinations and model inaccuracies can undermine decision quality, requiring robust guardrails, human-in-the-loop workflows, and domain-specific evaluation to validate outputs. Organizations also face trade-offs between



“Teams must integrate modern platforms with legacy systems, close critical talent gaps in AI and cybersecurity, and balance innovation with compliance.”

scaling AI quickly and controlling costs, while governing technology at scale through model risk management, observability, lineage/provenance tracking, and continuous monitoring—becomes essential for mission-critical applications where reliability, accountability, and rapid recovery are non-negotiable; cultural resistance further slows adoption even when the tech capability exists.

CIO&Leader: How has the cybersecurity conversation evolved from protection to enterprise-wide resilience, recovery, and accountability?

PREMKUMAR BALASUBRAMANIAN: Over the past few years, cybersecurity has moved from a perimeter-centric, “keep the bad guys out” mindset to an enterprise mandate for resilience, rapid recovery, and accountable risk management. We now operate on an assume-breach posture, architecting for continuity with zero-trust principles, real-time threat intelligence, and tested incident playbooks that compress recovery windows and protect revenue. Governance has matured as well: boards expect quantifiable risk reduction, clear ownership across business and technology, and

demonstrable compliance, with executives accountable for outcomes—not just controls. In short, security is no longer a defensive cost center; it’s a core capability that safeguards operations, brand, and customer trust at scale.

CIO&Leader: Which technology risks are now viewed as enterprise risks, and where are accountability boundaries still unclear between the CIO, CISO, and business leaders?

PREMKUMAR BALASUBRAMANIAN: Technology risks have firmly crossed into the enterprise risk domain, with boards now viewing cyber threats, data privacy breaches, AI misuse, and cloud concentration risks as strategic issues that can impact revenue, reputation, and regulatory standing. The challenge is that accountability boundaries remain blurred: who owns AI governance when models drive business decisions—the CIO for platforms, the CISO for controls, or business leaders for outcomes? Similarly, data integrity versus data security creates tension between CIO/CDO and CISO roles, while incident response often lacks clarity on final authority during crises. Even cloud risk management and third-party dependencies expose gaps, as procurement, security, and IT share overlapping responsibilities without a unified framework. The reality is clear: without explicit RACI models and board-level risk ownership, these gray zones will slow decision-making and amplify exposure as technology becomes inseparable from core business strategy.

CIO&Leader: How are boards reassessing cloud costs, ROI, and modernization priorities, and what tensions does this create for long-term digital strategy?

PREMKUMAR BALASUBRAMANIAN: Boards are scrutinizing run-

rate cloud spend vs. value, pushing FinOps discipline, rightsizing/repurchasing commitments, and prioritizing modernization that ties to revenue, resilience, and AI readiness (data platforms, app refactoring, security).

Tensions for long-term strategy:

- **Cost control vs. agility:** optimize spend without slowing innovation
- **Lift-and-shift vs. refactor:** quick savings vs. durable ROI and performance
- **Single-cloud simplicity vs. multi-cloud resilience:** vendor lock-in vs. portability
- **Centralized FinOps vs. decentralized product autonomy:** governance vs. speed
- **Near-term savings vs. AI/data investments:** opex reductions vs. future growth

CIO&Leader: What do CIOs need more from boards to deliver outcomes successfully, clearer direction, investment discipline, patience, or shared risk ownership?

PREMKUMAR BALASUBRAMANIAN: CIOs need clear strategic direction, consistent investment discipline, and above all, shared risk ownership to align technology bets with business priorities. Patience is critical for long-horizon initiatives like AI and modernization, but

“Security is no longer a defensive cost center; it’s a core capability that safeguards operations, brand, and customer trust at scale.”

boards must also actively champion change and accept trade-offs.

Top asks:

- **Clarity:** prioritize outcomes over projects
- **Discipline:** fund transformation, not just cost-cutting
- **Patience:** allow time for ROI on complex programs
- **Shared accountability:** cyber risk, data ethics, and innovation bets

CIO&Leader: What is the most underestimated technology risk or opportunity for 2026–27, and how do expectations differ across sectors?

PREMKUMAR BALASUBRAMANIAN: One of the most underestimated risks for 2026–27 is the integrity of AI models and the provenance of data. As organizations increasingly rely on generative and predictive AI, the rise of deepfakes, synthetic data, and unverified sources threatens trust and could trigger significant regulatory backlash. At the same time, the biggest opportunity lies in the convergence of AI and edge computing, enabling real-time automation in sectors like manufacturing, healthcare, and mobility—unlocking new revenue streams and operational resilience. Expectations vary widely across industries: financial services often underestimate compliance risks tied to AI while overestimating the near-term impact of quantum computing; healthcare sees transformative potential in AI diagnostics but faces challenges in patient data governance; manufacturing is leaning into autonomous operations but must contend with OT cybersecurity vulnerabilities; retail is betting on hyper-personalization while grappling with ethical AI and bias concerns; and the public sector views citizen services as a major opportunity but struggles with the pace of legacy modernization. ■



Gaurav Duggal
Senior Vice President - IT and Security at Jio Platforms

In five years, AI agents will be as essential as cloud infrastructure

Gaurav Duggal, Senior Vice President - IT and Security at Jio Platforms, explains how AI agents now manage core functions ranging from customer support in multiple languages to complex network operations

By **Musharrat Shahin** | editor@cioandleader.com

Reliance Jio is rapidly transforming how AI agents drive operations across its digital ecosystem. In a recent conversation with CIO&Leader, Gaurav Duggal, Senior Vice President – IT and Security at Jio Platforms Limited (JPL), shares insights into how India’s largest telecom and digital services company has moved beyond experimentation into full-scale deployment of AI agents, the business impact achieved so far, and the governance framework that ensures control, transparency, and accountability.

From customer support to network operations, fraud detection, and supply chain forecasting, Jio is redefining enterprise AI by balancing autonomy with rigorous oversight—a model that other large enterprises could emulate. Excerpts from the interaction.

CIO&Leader: How far along is Reliance Jio with AI agents, experimentation, pilots, or full production? Which functions are seeing real impact?

GAURAV DUGGAL: We are well beyond experi-

mentation. AI agents are operating in live production across several core functions. In customer support, they handle a large percentage of Level 1 interactions across voice and chat in multiple Indian languages. In network operations, agents monitor performance, detect anomalies, and trigger corrective actions within defined boundaries.

We are also using AI agents in fraud detection, marketing personalization, IT service management, and parts of supply chain forecasting. In enterprise sales and revenue assurance, we are running scaled pilots that are moving toward higher autonomy. The conversation internally is no longer about whether agents work. It is about how much autonomy we can responsibly and economically allow.

CIO&Leader: What is the most sophisticated AI agent you’ve deployed, and what measurable results has it achieved?

GAURAV DUGGAL: Our most advanced deployment is a closed loop network operations agent. It continuously monitors network telemetry,



“Our closed loop network operations agent has reduced incident resolution time by 30 to 40 percent and improved customer experience metrics such as latency and downtime.”

identifies anomalies, diagnoses probable root causes, simulates corrective scenarios, and executes approved configuration adjustments within guardrails.

The business impact has been significant. We have reduced incident resolution time by roughly 30 to 40 percent in certain clusters. We have lowered operational expenditure through fewer manual interventions and reduced field visits. We have also improved customer experience metrics such as latency and downtime, which directly correlates with churn reduction. The return on investment was achieved in under a year due to the scale of our network.

CIO&Leader: How do you determine which decisions AI agents can make on their own versus

requiring human oversight?

GAURAV DUGGAL: We classify decisions into three levels. First, fully autonomous decisions such as routine network parameter tuning within defined limits, Level 1 customer issue resolution, campaign optimization, and standard IT remediation.

Second, human-in-the-loop decisions such as pricing adjustments, high-value enterprise contract modifications, and certain fraud escalations where financial exposure crosses a threshold.

Third, strictly human-controlled decisions including capital allocation, regulatory submissions, major pricing strategy shifts, mergers and acquisitions, and brand-sensitive matters. The line is drawn based on financial exposure, regulatory impact, and reputational risk.

CIO&Leader: What level of operational or financial risk are you comfortable delegating to AI agents?

GAURAV DUGGAL: We are comfortable delegating low operational risk decisions and bounded financial risk decisions where predefined ceilings exist. Every autonomous workflow has transaction thresholds, loss caps, and escalation triggers embedded into the system.

We do not allow any AI agent to independently create a material balance sheet impact or take actions that could cause significant reputational harm. High regulatory and strategic risks remain fully human-governed. The principle is constrained autonomy, not blind automation.

CIO&Leader: Could you trace and explain the decisions made by AI agents if a regulator asked tomorrow?

GAURAV DUGGAL: Yes. Every production AI agent must meet strict auditability standards. We maintain full decision logs, model version traceability, policy rule mapping, and documented approval chains for guardrails.

If required by the Telecom Regulatory Authority of India or any other oversight authority, we can reconstruct the sequence from input data to model output to executed action. We can also demonstrate testing records for bias, robustness, and compliance. Explainability and governance are foundational in a regulated environment like ours.

CIO&Leader: As AI agents connect to core enterprise systems, how do you maintain control, visibility, and accountability?

GAURAV DUGGAL: Agents never receive unrestricted system access. We enforce role-based access controls, zero-trust architecture, and API-mediated access instead of



We classify decisions into three levels—fully autonomous, human-in-the-loop, and strictly human-controlled—based on financial exposure, regulatory impact, and reputational risk.

direct database writes. All actions are logged and monitored in real time.

We deploy sandbox simulations before allowing production-level execution. We also maintain kill switches and override capabilities. Integration is handled through a policy orchestration layer that defines what an agent can and cannot do. This ensures we do not lose control as autonomy increases.

CIO&Leader: Have you ever paused or shut down AI initiatives? What were the main challenges?

GAURAV DUGGAL: Yes, we have paused certain initiatives. In some cases, internal trust and change management broke before the technology did. Teams were not fully prepared for autonomous decision systems.

In other cases, we halted genera-

tive agents in sensitive workflows because hallucination risks were unacceptable. We have also rejected vendor models that lacked transparency and explainability. Governance and compliance concerns have been more limiting than technical capability.

CIO&Leader: What matters most in managing AI agents today: monitoring, human override, policy rules, or strict limits? How is control evolving?

GAURAV DUGGAL: Policy-defined boundaries matter most. Monitoring and human override are essential, but they are secondary to designing the right constraints upfront.

We are redefining control from approving every action to defining the permissible operating space within which agents function. Con-

trol becomes architectural rather than supervisory. The stronger the policy framework, the more confidently we can scale autonomy.

CIO&Leader: Looking ahead, which types of decisions will AI agents take over, and which will remain human-led? How will this evolve over three years?

GAURAV DUGGAL: In the next three years, AI agents will reliably own high-frequency, data-intensive decisions such as network micro-optimization, real-time fraud blocking, customer service resolution for standard issues, dynamic marketing personalization, and predictive inventory management.

Human judgment will remain essential for regulatory interpretation, ethical tradeoffs, crisis management, brand positioning, and long-term capital allocation. Machines will dominate operational velocity. Humans will focus on strategic direction and accountability.

CIO&Leader: Will AI agents become a central execution layer in enterprises, or will adoption be limited by regulation, architecture, or cost? Five-year outlook?

GAURAV DUGGAL: Five years out, AI agents will be a core execution layer within large enterprises like Reliance Industries Limited and its digital ecosystem. They will operate much like cloud infrastructure does today, embedded and indispensable.

However, scale will depend on governance maturity, regulatory clarity, secure infrastructure, and cost efficiency of models. The constraint will not be capability. It will be a leadership discipline in architecting responsible autonomy. Enterprises that institutionalize governance around intelligent systems will outperform those that treat AI as an isolated technology initiative. ■

Join **31,000+** CIOs Who Power the Most Engaged Tech Leadership Community on LinkedIn.

Join the exclusive **CIO&Leader LinkedIn Group**-a vibrant community where IT leaders like **YOU** come together to connect, collaborate, and share insights. With engagement levels higher than all our leading competitors combined, this is the ultimate platform to keep you informed, inspired, and ahead of the curve.

Discover curated content, leadership Insights, and thought leadership tailored for today's CIOs. Be part of the conversations that matter, learn from industry pioneers and network with the best minds in the industry.

The CIO&Leader community is your gateway to thought-provoking dialogue, cutting-edge tech & trends and actionable strategies.

Join the CIO&Leader LinkedIn Group today and elevate your leadership journey.

Follow us on **LinkedIn**
@CIOandLeader

Scan the QR Code to follow



You can also visit us at:
www.cioandleader.com

For more information, write to:
editor@cioandleader.com



CIO&LEADER

PRESENTS

CIO PLAYBOOK 2026

5th - 7th March 2026 * DoubleTree by Hilton, Varanasi

THANK YOU

for Making CIO Playbook 2026 a Defining Experience

TECHNOLOGY PARTNERS



TECHNOLOGY SHOWCASE PARTNERS



CYBERSECURITY RATING PARTNER



<https://events.cioandleader.com/>