JULY 2025 • VOLUME 14 • ISSUE 04 • ₹150

CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

AI: From Pilot to Production

Why AI Pilots Fail — and What India Inc Must Do to Scale

Pg. 12

→INSIDE: Scalable AI frameworks, industry-specific CIO strategies, and leadership lessons from India's digital frontrunners.

A 9.9 GROUP PUBLICATION
cioandleader.com
cioandleader cioandleader/

CIO&LEADER studiotalks

CIO&LEADER STUDIOTALKS—
WHERE TECHNOLOGY MEETS THE SPOTLIGHT!

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

WHY JOIN STUDIOTALKS?

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

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

Be featured among India's top technology leaders.

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

SECURE YOUR SPOT NOW!

For more information
Jatinder Singh
Executive Editor - Enterprise Tech
jatinder.singh@9dot9.in, +919718154231

For Business Proposal
Hafeez Shaikh
National Sales Head, B2B Tech,
hafeez.shaikh@9dot9.in, +91 9833103611

Follow us: @CIOandLeader

Wanted: More scalable AI use cases

AS WE began planning the 26th Annual CIO&Leader Conference, scheduled for August 1–3 in Pune, one question consistently dominated our discussions: What is truly top of mind for today's CIOs?

While "The ROI of AI" emerged as a strong contender, our conversations with CIOs and analysts revealed a deeper, more pressing concern: the ongoing struggle to move AI initiatives from pilot to production. Despite growing boardroom pressure to demonstrate tangible business value from AI, many organizations remain stuck in the experimentation phase—unable to scale promising proofs of concept into enterprise-grade deployments.

This hesitation isn't driven by a lack of intent, but rather a lack of scalable, repeatable use cases that deliver measurable outcomes. AI is no longer a fringe experiment; it is fast becoming a central pillar of enterprise transformation. Today, technology is inseparable from business strategy—and AI, along with emerging autonomous systems, sits at the heart of that evolution.

According to a recent BCG report, successful AI adoption depends on three imperatives: mitigating and communicating risks, building an enterprise-wide foundation for scale, and establishing clear policies on tool usage, governance, and compliance. None of this is simple. But CIOs—positioned at the intersection of business vision and technological capability—are uniquely equipped to lead.

As the conversations at the upcoming conference will explore, the next phase of AI isn't about models or algorithms—it's about operationalization. And that requires more than innovation; it demands leadership.

That's why this issue—and our cover story—sets the stage for the conversation ahead. At our Annual Conference, top CIOs, analysts, and business leaders will gather to explore how to operationalize AI at scale and unlock its full enterprise value. The next wave of AI isn't about innovation alone, it's about execution. From pilots to production. From ideas to impact. ■

"This issue—and its cover story—previews our CIO&L Annual conference in August, where top CIOs, analysts, and business leaders will explore how to scale AI from pilot to production and unlock its full enterprise value."

Jatinder Singh
Executive Editor
jatinder.singh@9dot9.in

content July 2025

COVER STORY

Cover Design by:
Manish Kumar

Please Recycle This Magazine And Remove Inserts Before Recycling

NEWS & VIEWS

AI: FROM PILOT TO PRODUCTION

INSIGHT

TECH TALK

CIO&LEADER

www.cioandleader.com

MANAGEMENT
Managing Director: **Dr Pramath Raj Sinha**
Printer & Publisher / CEO & Editorial Director (B2B Tech): **Vikas Gupta**
COO & Associate Publisher (B2B Tech): **Sachin Nandkishor Mhashilkar (+91 99203 48755)**

EDITORIAL
Group Editor: **R Giridhar**
Executive Editor: **Jatinder Singh**
Principal Correspondent & Editorial Coordinator – CIO&Leader: **Musharrat Shahin**
Correspondent: **Jagrati Rakheja**

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

SALES & MARKETING
Director - B2B Tech:
Vandana Chauhan (+91 99589 84581)
National Sales Head - B2B Tech:
Hafeez Shaikh (+91 98331 03611)
Head - Brand & Strategy:
Rajiv Pathak (+91 8010757100)

COMMUNITY ENGAGEMENT & DEVELOPMENT
Head - Community Relations: **Dipanjan Mitra**
Head - Databases: **Neelam Adhangale**
Community Manager: **Vaishali Banerjee**
Community Manager: **Snehal Thosar**
Community Manager: **Reetu Pande**
Community Manager: **Nitika Karyet**
Assistant Manager Community Development: **Shabana Shariff**

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

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

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 Tara Art Printers Pvt Ltd., A-46-47, Sector-5, NOIDA (U.P.) 201301.

Editor: **Vikas Gupta**

9.9 GROUP

CIOmovements

Sharad Garg elevated to Vice President – Services at SymphonyAI

Sharad Garg is promoted to Vice President – Services at SymphonyAI, leading project delivery and client success across US and Europe with 17+ years' experience.

Sunil Pandey steps into the Role of Chief AI and Digital Officer at HFCL Limited

Sunil Pandey has been appointed Chief AI and Digital Officer at HFCL, driving AI-led innovation and digital acceleration across business verticals with two decades' expertise.

Sudeep Agrawal joins Reliance Infrastructure as Chief Information Officer

Sudeep Agrawal becomes CIO at Reliance Infrastructure, bringing two decades of expertise to lead IT strategy, innovation, and digital transformation across the enterprise.

Raghavendra Vaidya appointed as Group CIO, Daimler Truck AG & Board Member, Daimler Truck Innovation Center India

Raghavendra Vaidya becomes Group CIO at Daimler Truck AG and Board Member at DTICI, driving global digital transformation with two decades of enterprise IT leadership.

Sanjay Sehgal elevated to CEO & Managing Director at TP-Link India

Sanjay Sehgal is promoted to CEO and MD at TP-Link India, leading innovation and market expansion with two decades of expertise in networking and enterprise solutions.

Manish Gupta appointed as President and Managing Director – India at Dell Technologies

Manish Gupta is appointed President and MD of Dell Technologies India, bringing 20+ years' leadership experience to drive innovation, growth, and digital transformation across the region.

Jitendra Mangave joins Metro Brands Limited as CIO & CTO

Jitendra Mangave becomes CIO & CTO at Metro Brands, leading digital innovation with 20+ years' experience in IT strategy, transformation, and enterprise retail technology leadership.

Ravi Jayanthi joins Nestavia Home Finance as Chief Technology Officer!

Ravi Jayanthi becomes CTO at Nestavia Home Finance, bringing two decades of expertise to lead technology strategy, innovation, and IT operations for tech-driven financial solutions.

Sanjay Verma joins Easa Husain Alyousifi & Sons Co. as Director of IT

Sanjay Verma becomes Director of IT at Easa Husain Alyousifi & Sons, bringing 30+ years' experience to drive digital transformation and infrastructure modernization across the enterprise.

Binita Prasad takes charge as Head of IT & Digital at Schindler India

Binita Prasad becomes Head of IT & Digital at Schindler India, bringing 30+ years of expertise to lead digital strategy, innovation, and business-aligned technology transformation.

NEXT100 2017 winner Amit Kumar joins Manipal Health Enterprises as Director – IT

NEXT100 2017 winner Amit Kumar becomes Director – IT at Manipal Health, bringing 20+ years' expertise to drive digital innovation and technology transformation in healthcare.

Pocket Entertainment elevates Umesh Bude to CTO

Umesh Bude is elevated to CTO at Pocket Entertainment, leading AI-powered storytelling and tech innovation across Pocket FM, Pocket Toons, and Pocket Novel platforms.

Veteran Tech Leader Manish Israni Appointed CTO – India & SAARC at Westcon-Comstor

Manish Israni becomes CTO – India & SAARC at Westcon-Comstor, leading cloud, cybersecurity, and digital strategy with two decades of enterprise IT and infrastructure experience.

Vertiv Names Mike Giresi as Global CIO

Mike Giresi joins Vertiv as Global CIO to lead AI, cybersecurity, and digital strategy, enhancing customer experience, operational leverage, and executional strength globally.

Sourav Dasgupta Appointed Chief Information Officer at Allcargo

Sourav Dasgupta is promoted to CIO at Allcargo, bringing 20+ years of expertise to lead technology vision and digital transformation for tech-enabled growth and excellence

Manoj Kumar Mauni joins Airtel as CIO – Telecom

Manoj Kumar Mauni becomes CIO – Telecom at Airtel, bringing 30 years' experience to lead IT strategy, service delivery, and digital innovation across the organization.

Rajgopal Nayak joins IndoSpace as Executive Vice President – Information Technology

Rajgopal Nayak becomes EVP – IT at IndoSpace, bringing 20+ years of expertise to lead technology vision, digital transformation, and innovation in industrial real estate.

Sanjay Nandavadekar appointed as Head of IT at Inventia

Sanjay Nandavadekar becomes Head of IT at Inventia, bringing 20+ years of pharma IT leadership to drive digital transformation, compliance, and innovation in regulated environments.

Anand Tomar joins Allied Blenders & Distillers Ltd. as Chief Information Officer

Anand Tomar becomes CIO at Allied Blenders & Distillers, bringing 20 years' experience to lead technology strategy, cybersecurity, and digital transformation from the Mumbai headquarters.

Mohan Shah joins ClaimPro Assist Pvt. Ltd as Chief Technology Officer

With nearly three decades of extensive experience in IT leadership, digital transformation, and infrastructure management, Mohan will lead the company's technology strategy, driving innovation and operational excellence .

AI Revolution Reshapes Software Development

Gartner predicts AI will transform software development, with 90% of engineers using AI tools and LLMs shaping intelligent applications.

By **Jagrati Rakheja** | jagrati.rakheja@9dot9.in

THE ERA of human-dominated coding is rapidly giving way to AI-powered development, according to research firm Gartner's latest strategic outlook for software engineering.
Released Wednesday, Gartner's report identifies six critical trends that will define how software gets built over the next three years, with artificial intelligence serving as the primary catalyst for change.

AI Takes Center Stage in Development

The most striking prediction centers on AI adoption rates. Gartner forecasts that 90% of enterprise software engineers will use AI coding assistants by 2028—a dramatic surge from just 14% in early 2024.

"AI-enabled tools and technologies are fundamentally changing how software is built and delivered," said Joachim Herschmann, VP Analyst at Gartner. This shift will transform developers from hands-on coders into orchestrators who focus on problem-solving and system design.

AI is turning developers into system designers—shifting focus from writing code to orchestrating intelligent, sustainable solutions.

Building Intelligent Applications

Large language models are becoming integral to software development. By 2027, at least 55% of engineering teams will actively build applications featuring LLM-based capabilities, enabling software to interact more like humans.

Platform Engineering Gets AI Boost

Organizations are integrating AI capabilities into their internal development platforms. Gartner predicts 70% of companies with platform teams will embed generative AI features by 2027, making AI tools easily accessible through self-service portals.

Focus on Talent and Sustainability

Two additional trends highlight evolving priorities: maximizing "talent density" by concentrating highly skilled professionals within teams, and embracing open AI models that offer greater flexibility and lower costs.

Environmental concerns also drive the push toward "green software engineering," incorporating carbon-efficient practices from project planning through production.

The trends collectively point toward a future where AI augments human creativity while organizations balance innovation with sustainability and talent optimization.■

Scindia Pitches 6GHz Delicensing as Key to India's Digital Future

India will delicense the 6GHz band by August 15, 2025, unlocking high-speed, inclusive Wi-Fi access for all citizens.

By **Jagrati Rakheja** | jagrati.rakheja@9dot9.in

IN A landmark announcement on World Wi-Fi Day, Union Minister for Communications Jyotiraditya M. Scindia said that India will delicense the 6GHz spectrum band by August 15, 2025, marking a major push towards universal digital access and high-speed public Wi-Fi, especially in underserved and rural regions.

This is not just a policy change. It's a shift in mindset. We are no longer just regulators—we are facilitators," said Scindia at a conference organised by the Broadband India Forum (BIF) in the capital. "Connectivity is not a commodity; it is an act of nation-building," he added, urging stakeholders to innovate around affordability, competition, and last-mile access.

Wi-Fi for All: A Vision of Inclusive Growth

Delicensing the 6GHz band will power next-gen Wi-Fi like 6E and 7, enabling faster speeds, lower latency, and seamless experiences for 8K streaming, AR, and remote healthcare. Minister Scindia highlighted Wi-Fi's transformative role in India's digital growth, projecting a $22 billion market by 2035. "Every hotspot must become a hope spot," he said, noting India's global lead in affordable data—just ₹9 per GB compared to the $2.49 global average.

Addressing the Digital Divide with PM-WANI

Former TRAI Chairman R.S. Sharma called PM-WANI a "UPI-like revolution" in telecom. Designed as a decentralised public Wi-Fi framework, PM-WANI aims to democratise broadband by empowering small entrepreneurs to create Wi-Fi hotspots. Scindia highlighted ongoing pilots across 13 villages testing 5G-based services in education, healthcare, and agriculture.

No to Duopoly, Yes to Healthy Competition

Taking aim at market concentration in telecom, Scindia said, "Duopoly is not good. We must have competition in every vertical." India's telecom sector is currently dominated by Reliance Jio and Bharti Airtel, with other players like Vodafone Idea and BSNL struggling to scale 4G and 5G services. The minister also hinted that spectrum for satellite internet services will be allocated soon on an administrative basis.

Why 6GHz Matters Now

According to BIF Chairperson Aruna Sundararajan, more than 80% of India's internet usage happens indoors, making robust Wi-Fi essential for digital inclusion. "Delicensing the 6GHz band is the key to leapfrogging into next-generation connectivity," she said, adding that mobile networks alone cannot meet the rising data demand.

A Policy Built on Empowerment, Not Exclusivity

In conclusion, Scindia framed the delicensing of the 6GHz band as a symbol of India's commitment to technology neutrality, customer-centricity, and inclusive growth. ■

One in Ten GenAI Apps Pose Serious Security Risks, Warns Report

Palo Alto Networks warns that 1 in 10 GenAI apps used by enterprises pose high security risks amid soaring AI adoption.

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

AS GENERATIVE AI becomes deeply embedded in enterprise workflows, a new report by Palo Alto Networks reveals a concerning reality: 10% of the average company's GenAI apps are considered high-risk.

What if the very AI tools driving workplace productivity were also opening the backdoor to cyber threats? That's the concern raised by Palo Alto Networks in its latest State of Generative AI 2025 report. The cybersecurity giant reveals that on average, 10% of the 66 GenAI apps used in a typical enterprise today fall under the high-risk category.

Based on data from over 7,000 global enterprises, the report underscores a critical contradiction: as Generative AI (GenAI) adoption accelerates, security controls have not kept pace—particularly across Asia-Pacific and India.

Rapid Growth, Rising Risk

The findings show an 890% spike in GenAI traffic in 2024, with India emerging as a major adopter. Popular apps like Grammarly, Microsoft Power Apps, and Copilot dominate usage across Indian enterprises. But beneath the surface lies growing concern over Shadow AI—unauthorized or unmonitored AI use—making it hard for IT teams to track data exposure and enforce compliance.

"AI is transforming how governments and businesses work," says Tom Scully, Director at Palo Alto Networks, Asia Pacific & Japan. "But that innovation must be matched with strong oversight. Otherwise, the very tools we celebrate could compromise national security, public trust, and business integrity."

The very tools driving innovation could also compromise national security, public trust, and business integrity.

GenAI Use Cases... and Misuse

The launch of DeepSeek-R1 in January 2025 triggered a 1,800% spike in related traffic within two months.

- Data loss prevention (DLP) incidents tied to GenAI more than doubled in 2025, now accounting for 14% of all security incidents.
- AI models remain vulnerable to jailbreak attacks, producing harmful or unsafe outputs.

Securing the Future of AI

Palo Alto Networks recommends three key actions for businesses:

- Establish visibility and control of GenAI use across the organisation.
- Safeguard sensitive data with real-time monitoring and policy enforcement.
- Adopt Zero Trust architectures to defend against evolving AI-powered attacks. ■

AI Now Writes Over Half of All Spam Emails, Finds New Study

Over 51% of all spam emails are now AI-generated, as cybercriminals use artificial intelligence to scale and refine attacks.

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

RESEARCHERS REVEAL a sharp rise in AI-generated spam, as cybercriminals use artificial intelligence to automate and refine email attacks.

A New Kind of Inbox Invader

Check your spam folder—over half of it is now AI-generated. A new study by Columbia, the University of Chicago, and Barracuda Networks found that 51% of spam emails are created by AI tools, a sharp increase since ChatGPT's public launch in November 2022.

Researchers from Columbia, the University of Chicago, and Barracuda Networks found that AI-generated spam now makes up 51% of all spam emails—a sharp increase since ChatGPT's launch in November 2022. The study analyzed spam and malicious emails from February 2022 to April 2025.

Spam Gets Smarter, Not Just Louder

The study shows a steady growth in AI use across email threats:

- 51% of spam emails in April 2025 were AI-generated
- 14% of business email compromise (BEC) attacks used AI—though still mostly human-written
- AI-generated messages tend to be more formal, grammatically accurate, and cleverly phrased
- Cybercriminals are using AI to tweak wording in spam emails to evade security filters. Researchers note that while AI isn't changing the nature of attacks, it's improving how they're written.

"By April 2025, the majority of spam emails were not written by humans," said Prof. Asaf Cidon of Columbia University. "While targeted attacks like BEC still require a human touch, AI's role is steadily growing."

A Challenge to Detect and Defend

To determine AI's role in email creation, researchers compared messages from before and after ChatGPT's launch, training systems to detect AI-written content. Still, spotting AI involvement remains a challenge. "We can only see the output, not how it was created," added Cidon.

India Not Immune to the Threat

Parag Khurana, Country Manager, India at Barracuda Networks, stressed the rising threat: "Cybercriminals are using AI to scale attacks. It's crucial for Indian organisations to adopt platform-based defenses that offer deep threat visibility and rapid response."

Barracuda recommends combining AI-powered email protection with cybersecurity training to guard against increasingly sophisticated spam and phishing attempts. ■

80% of GenAI Apps Will Use Existing Data – Is Yours Ready?

Gartner forecasts that 80% of GenAI business apps will be built on existing data platforms by 2028, making RAG a critical enabler of faster, safer AI deployment.

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

IN A world racing to embrace Generative AI (GenAI), the question isn't whether businesses will adopt it, but how fast and how wisely. At the Gartner Data & Analytics Summit in Mumbai, the conversation has shifted from hype to how-to. Gartner predicts that by 2028, 80% of GenAI business applications will be developed on existing data management platforms, cutting development time and complexity by half.

This trend marks a pivotal shift toward using in-house data systems as a launchpad for GenAI innovation, reducing the cost and chaos of integrating fragmented tools and frameworks.

Why RAG Matters Now

A key theme at the summit is the growing adoption of Retrieval-Augmented Generation (RAG)—a technique that combines the power of large language models (LLMs) with real-time access to business-specific data. According to Prasad Pore, Senior Director Analyst at Gartner, RAG offers "implementation flexibility, enhanced explainability, and composability" for GenAI apps.

Since most LLMs are trained on public data, they often fail to deliver when solving organization-specific problems. RAG bridges this gap by feeding the LLMs with internal data, structured and unstructured, while using tools like vector search, metadata tagging, and chunking to enhance accuracy and trust.

Recommendations for Enterprises

To fully leverage GenAI, Gartner advises businesses to:

- Transform data platforms into RAG-ready systems that can act as real-time knowledge sources.
- Adopt RAG technologies like vector search and graph databases, ensuring resilience and scalability.
- Use metadata strategically, technical and operational, to secure GenAI from privacy risks and misuse.

As India's GenAI ecosystem accelerates, highlighted by recent moves from major tech players like Infosys and TCS to integrate GenAI into core services, RAG is emerging as a vital foundation, not just an add-on. ■

How Gen AI is Redefining ROI in the Enterprise Era

In a world saturated with data but starved for actionable insight, Generative AI is emerging as a decisive force in how enterprises measure and deliver business value.

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

A NEW whitepaper released by business consulting firm Practus and AI transformation company Pathsetter AI, titled "The Future of Enterprise Intelligence: Integrating GenAI for Competitive Advantage," argues that traditional ROI models are being upended by Gen AI's capacity to deliver real-time, predictive, and outcome-oriented decision-making.

For decades, ROI focused primarily on efficiency metrics—cost reduction, productivity gains, and historical performance. However, with 97% of data leaders admitting they struggle to extract meaningful value from their data, the cracks in conventional Business Intelligence (BI) systems are showing. Static dashboards, delayed reports, and manual data prep are no longer sufficient in a hyper-competitive, real-time economy.

Gen AI, by contrast, offers a paradigm shift. It enables enterprises to anticipate trends, automate decision cycles, and personalize customer engagement at scale.

The whitepaper showcases transformative results: Klarna's AI-powered customer support ecosystem resolved 67% of queries autonomously and slashed response times by 78%, while Coca-Cola's Gen AI content engine cut production time by 30% and boosted localized engagement by 15%. These aren't isolated tech pilots—they're measurable business wins.

A standout concept in the report is Outcomes as a Service (OaaS), a model where Gen AI projects are judged not by features deployed, but by KPIs like reduced churn, improved forecasting accuracy, or accelerated time-to-decision. For example, a Gen AI engagement is considered successful not for launching a forecasting tool but for achieving "10% improvement in revenue forecast accuracy."

The shift also calls for a more nuanced set of KPIs. The paper proposes a new ROI lexicon that includes metrics like Innovation Rate (target: 25–40%), Forecast Accuracy (85–95%), and AI Adoption Rate (30–70% annually). This moves enterprise performance from merely tracking outputs to orchestrating outcomes.

Yet challenges persist. Legacy IT systems, cultural resistance to automation, and concerns around AI bias and data governance remain significant roadblocks. The paper cautions against viewing Gen AI as a plug-and-play solution. Execution demands internal alignment, iterative model training, and human-AI collaboration.

Still, the direction is clear. Enterprises that embed Gen AI not just in tools but in their operating models are likely to outperform those who hesitate.

As Gen AI moves from pilot projects to performance engines, organizations must realign their strategic lens. ROI is no longer a measure of what was saved—it's a map of what's possible. ■

AI: From Pilot to Production

Why AI Pilots Fail – and What India Inc Must Do to Scale

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

AI is no longer an emerging technology—it's a boardroom mandate. Across industries, enterprise leaders are under growing pressure to demonstrate results, not just run pilots.
By the end of 2024, more than 78% of Indian enterprises had implemented AI in at least one function, reflecting rapid adoption, according to McKinsey.

"People didn't believe in AI until they saw it. But once they did, the conversation changed from 'can it do this?' to 'can it do more? Now they ask, 'Can it do more?'" says Vinod Khode, Group CIO at Varroc Engineering.

However, a global study by BCG reveals that 74% of organizations continue to face significant challenges in scaling AI beyond initial pilots. This highlights a critical gap—not in willingness to adopt AI, but in the ability to operationalize it across business functions.

For Indian CIOs, the challenge is not about experimenting with AI—it's about scaling it across complex, legacy-laden environments, aligning it with business goals, and ensuring it's both responsible and resilient. The real test lies in moving AI from innovation labs and test beds into the operational core of the enterprise.

For enterprise CIOs, the transition from experimentation to enterprise-wide execution remains one of the most complex and strategic challenges, requiring alignment of people, processes, data, and infrastructure to deliver sustained business value. However, research from Accenture highlights that organizations that have managed to scale AI effectively can outperform their peers by as much as 2.5X in both revenue growth and productivity.

We look at how leading CIOs in India are addressing these challenges. From manufacturing and digital commerce to automotive and engineering, these leaders are reimagining AI as an integrated, measurable, and secure business enabler. Their experiences provide a practical roadmap for enterprise technology leaders who are ready to move beyond experimentation and deliver enterprise-scale impact.

Moving from Ideas to Impact

Successful AI transformation begins with strategic intent, not pilot enthusiasm.

Turning AI pilots into enterprise-scale success requires more than technical capability—it demands clarity of purpose. Moving from ideas to impact begins with defining what success looks like and ensuring every use case is tied to business value. Strategic intent—not experimentation—must lead the way.

CIOs are realizing that without top leadership alignment and measurable outcomes, AI risks remaining a lab experiment. As organizations mature, they focus on achieving early wins, building executive buy-in, and preparing their teams for what comes next: designing the proper infrastructure to scale AI confidently, reliably, and with continuity across all functions.

Vinod Khode started with executive education before moving on to deployment. He invested early in change management, ensuring cultural readiness across teams.

Each use case should be mapped to specific business outcomes, such as cost reduction and increased productivity. "If it's not aligned to business, it won't earn trust—or budget," Vinod Bhat of Tata AuoComp explains. "You don't walk into the AI Olympics. You prepare for it, years in advance," advises Bhat

Like many Indian enterprises, Varroc's AI journey began with cautious curiosity. "We started with educating a focused core team," Khode explains. Rather than rushing into implementation, the team started with focused

"People didn't believe in AI until they saw it"

—Vinod Khode
Group CIO, Varroc Engineering.

"You need model explainability not just for trust, but for regulators, partners, and your teams,"

—Vinod Bhat
Chief Digital Officer,
Tata AutoComp.

group discussions, supplemented by consultations with Gartner to learn how similar organizations worldwide are applying AI, from operations and quality to HR.

A clear strategy and organizational readiness are essential to move AI beyond isolated pilots. When initiatives are aligned with business outcomes and supported by informed leadership, they are more likely to progress toward scale. The next step is building a foundation that enables reliable, secure, and scalable AI deployment across the enterprise.

Building the Foundation

AI requires a scalable and flexible architecture. After setting clear objectives, the next step in scaling AI is establishing a robust and flexible technology foundation. Scalable AI requires infrastructure that can support real-time data processing, seamless integration, and the ability to evolve with emerging technologies.

Modern AI workloads, especially those involving deep learning, rely heavily on high-performance hardware. GPUs (Graphics Processing Units) have become the industry standard for training complex AI models due to their parallel processing capabilities, which significantly reduce the time required for model development. On the other hand, CPUs (Central Processing Units), while slower for training, remain crucial for lightweight inference tasks and general-purpose computing. As enterprises move from training to deployment, many are adopting hybrid compute environments—leveraging GPUs for training and CPUs or optimized inference accelerators (like TPUs or FPGAs) for running models at scale across diverse workloads and environments. This strategic use of hardware ensures cost-efficiency without compromising performance.

Indian enterprises are prioritizing architectures that strike a balance between performance and adaptability, leveraging cloud-native environments, edge computing, and containerized services to enable agile deployment and scalability. These foundational choices not only accelerate the transition from pilot to production, but also prepare the organization for long-term growth. Once the infrastructure is in place, the focus shifts to managing the data layer effectively to support intelligent, scalable insights.

Indian enterprises are prioritizing architectures that strike a balance between performance and adaptability, leveraging cloud-native environments, edge computing, and containerized services to enable agile deployment and scalability. These foundational choices not only accelerate the transition from pilot to production, but also prepare the organization for long-term growth. Once the infrastructure is in place, the focus shifts to managing the data layer effectively to support intelligent, scalable insights.

At IndiaMART, CIO Nikhil Prabhakar has led the development of a future-ready technology architecture designed to support the evolving demands of AI at scale. The company has implemented a microservices-based framework, powered by Docker and Kubernetes, that operates within a resilient multi-cloud environment.

This approach enables greater flexibility, faster deployment cycles, and simplified maintenance. "Our architecture is modular and open-ended," Prabhakar explains. "It's designed in a way that allows us to integrate emerging AI innovations seamlessly, without having to overhaul or disrupt our core systems." This foundation not only supports rapid experimentation but also ensures long-term adaptability as AI technologies mature.

At Force Motors, Group CIO Anand Deodhar has overseen the implementation of a hybrid edge-cloud architecture designed to meet

the performance demands of AI-driven manufacturing. The company has deployed GPU-powered workstations at the edge to enable near real-time inference, reducing reliance on centralized cloud processing for time-sensitive tasks.

"We've built a hybrid edge-cloud architecture to support low-latency AI in production," Deodhar explains. "From predictive maintenance to dynamic shift planning—latency matters." This infrastructure enables AI applications to respond quickly to operational data, thereby enhancing decision-making on the factory floor and facilitating more efficient, intelligent production workflows.

A well-designed infrastructure enables AI systems to scale efficiently and adapt to changing business needs. With a flexible, modular, and performance-driven architecture in place, organizations are better equipped to support a diverse range of AI applications. The next phase involves focusing on data, ensuring it is reliable, accessible, and ready to drive intelligent outcomes.

Fueling Scalable Insights

Data readiness is a prerequisite. For AI to move beyond isolated success stories, enterprises must focus on data readiness, ensuring the right volume, quality, and governance of data is in place. This approach balances innovation with cost-efficiency and compliance.

Tata AutoComp has adopted a unified data strategy, incorporating real-time model monitoring and retraining frameworks, to ensure accurate and relevant outputs. "From clean pipelines to explainability audits—we've baked governance into the process," notes Vinod Bhat.

At Varroc Engineering, Vinod Khode followed a leaner model by repurposing existing datasets and introducing automated labeling, enabling scale with control. As data systems mature, the focus naturally shifts to ensuring that AI insights can empower users across functions.

"Before acquiring new data, we repurposed internal datasets. Structured data governance and automated labeling followed," says Vinod Khode.

A strong data foundation ensures AI models remain relevant, transparent, and effective over time. With governance, reusability, and retraining mechanisms in place, enterprises can scale insights reliably. The next priority is to integrate these insights into daily operations, ensuring that AI tools are accessible, practical, and applicable across all business teams.

AI: From Shop Floors to Startups

Indian enterprises are deploying AI where it matters:

- **Manufacturing:** Dixon, Varroc, and Force Motors use AI for predictive maintenance, defect detection, and production planning.
- **B2B Commerce:** IndiaMART uses AI for fraud detection, semantic enrichment, and automated buyer-seller negotiation bots.
- **Automotive:** Maruti Suzuki's AI-powered mobile manual utilizes image recognition to support customers in real-time.

Making AI Everyone's Business

For AI to create value, it must embed into daily workflows. Once the infrastructure is in place, the next critical step is enabling AI systems with high-quality, accessible, and well-governed data. Scalable AI insights rely on continuous data flow, model retraining, and real-time feedback mechanisms to ensure relevance and accuracy.

"We've built a hybrid edge-cloud architecture to support low-latency AI in production,"

—Anand Deodhar
Group CIO, Force Motors

CIOs are prioritizing unified data strategies, centralized governance, and clean pipelines to support sustained model performance across use cases. These practices enable AI systems to evolve in response to changing business needs, providing a reliable foundation for deployment. As organizations harness the full potential of data-driven intelligence, the next challenge lies in embedding AI into the daily decision-making fabric across functions and teams.

At Dixon Technologies, AI is integrated across various functions, including quality checks, inventory planning, and training. "Our AI-powered planning engine forecasts procurement needs, reducing material waste," says Pradhan. "We've also deployed AI dashboards that empower shop floor teams to make real-time, data-backed decisions."

Force Motors uses AI to inform shift planning and machine utilization. "We don't treat AI as a separate project. It's embedded into our plant DNA," says Deodhar. "In manufacturing, success is when the machine learns—and then teaches the team."

With structured, scalable data frameworks in place, enterprises will be better positioned to generate consistent and actionable AI insights. Ensuring model performance through ongoing retraining and clean data pipelines is essential. The next phase is integrating these insights into routine workflows—making AI accessible, usable, and valuable across business functions.

"Having AI and succeeding with AI are two entirely different challenges,"

—Amit Pradhan
VP & CIO,
Dixon Technologies

Measuring What Matters

According to Deloitte, most enterprises take over a year to achieve ROI from AI. The CIOs profiled here show that the right metrics can accelerate this curve. As AI moves into production, enterprises are shifting focus from experimental outputs to real-world outcomes—whether it's improved customer experience, operational efficiency, or cost optimization.

IndiaMART, for instance, uses KPIs such as responsiveness, innovation impact, and security enhancements to evaluate AI initiatives. Force Motors links AI outcomes directly to production performance indicators, such as uptime and energy efficiency. With the right metrics in place, enterprises can accelerate ROI and scale with confidence

IndiaMART's Prabhakar shares that their WhatsApp-based IM Insta bot tripled responsiveness, while PhotoSearch 3.0 reduced search latency across millions of products. "Every AI initiative is evaluated not by hype, but by impact: CX improvement, innovation, and security. GenAI and Agentic AI are not just ideas for us—they're delivering measurable ROI," says Prabhakar

Force Motors tracks AI's contribution to uptime, throughput, and energy usage through machine performance scorecards. As AI becomes embedded in business operations, the next priority is to ensure that systems are not only effective but also trusted.

Measuring AI initiatives against relevant business KPIs enables organizations to track value, enhance decision-making, and support broader adoption of AI. When performance indicators reflect real impact, AI can scale with purpose. The next focus is to strengthen trust through transparency, governance, and responsible deployment across all layers of the enterprise.

Building Trust into AI

WIth AI becoming deeply embedded into business operations, trust is emerging as a non-negotiable foundation for scale. According to IBM, 85% of executives believe trustworthy AI is essential for business success. CIOs are responding by implementing robust gover-

nance frameworks to prioritize transparency, fairness, and security from the outset.

Varroc Engineering's early AI pilots were grounded in solving practical business problems, helping build trust in the technology across teams. From a conversational BI tool that empowers leaders with instant insights to a 24/7 virtual HR assistant, and an AI guide for onboarding through PLM systems, each use case delivered clear, reliable outcomes. As employees experienced tangible benefits, confidence in AI grew organically, laying the foundation for broader enterprise adoption.

At IndiaMART, building trust in AI is a core priority. They have established a robust internal AI governance framework to ensure our systems are ethical, explainable, and free from bias or hallucinations, especially as they integrate Generative AI into key functions, such as lead verification and communication. IndiaMART is also advancing the convergence of AI and cybersecurity, using intelligent models that proactively detect threats, prevent fraud, and mitigate abuse before it affects users.

At Maruti Suzuki, the foundation of effective and trustworthy AI lies in clean, structured, and well-labeled data. Whether leveraging small or large language models, the focus remains on aligning high-quality data with clear business objectives. Each AI deployment is guided by robust governance frameworks designed to mitigate risks such as hallucination and bias. By prioritizing strong data management and embedding safeguards at every stage of development, the organization aims to ensure that AI systems are not only intelligent but also reliable, responsible, and capable of delivering real value across the enterprise.

."It's how we make our enterprise more human, efficient, and future-ready. Customers don't want tools. They want intelligent solutions," says Dr. Tapan Sahoo, Head of Digital Enterprise, Information & Cyber Security, Maruti Suzuki India.

AI That Scales Is AI That Succeeds

While AI experimentation is widespread, actual business value lies in scale, not pilots. According to a recent NetApp survey, 81% of global organizations are currently piloting or scaling AI initiatives, but only 4% have successfully industrialized AI across multiple business

"Every AI initiative is evaluated not by hype, but by impact,"

—Nikhil Prabhakar
CIO, IndiaMart.

Key Takeaways

- **Strategy First, Technology Follows:** Every successful AI initiative begins with alignment to business goals. Without a clear use-case-to-impact map, scale will stall.
- **Build Infrastructure That Enables Agility:** Containerized, cloud-native, and modular architectures enable rapid experimentation and integration without disruption.
- **Invest in Data Governance and Quality:** Clean, trusted data and retrainable pipelines are critical to avoiding model decay and ensuring AI performance at scale.
- **Integrate AI into Daily Workflows:** AI must empower business users, not just data scientists. Democratize access through intuitive dashboards and user-centric design.
- **Measure What Matters:** Tie AI projects to specific KPIs. Use scorecards, business reviews, and operational metrics to validate ROI.
- **Drive Culture, Not Just Code:** Upskill teams, create internal champions, and nurture innovation communities. Cultural readiness is a prerequisite for adoption.
- **Governance Is Non-Negotiable:** Build transparency, fairness, explainability, and compliance into AI systems from day one—especially in regulated industries.
- **Think Platforms, Not Pilots:** Use platform thinking to build AI capabilities that can scale horizontally across departments and functions.

The Maturity Journey: A CIO's Guide to Scaling AI				
Stage	**Focus Area**	**Challenges Faced**	**Key Success Indicators**	**CIO Insights & Examples**
1. Exploration & Pilots	Idea validation, isolated use cases, internal demos	Limited business alignment, fragmented data, lack of AI literacy	Working POC, stakeholder curiosity, initial buy-in	Dixon: 20+ POCs across manufacturing; Verrock: Virtual HR, NLP BI
2. Early Adoption	Scaling select use cases within a function, building architecture foundation	Siloed pilots, lack of MLOps, integration with legacy systems	Departmental impact, team enthusiasm, early ROI	IndiaMART: IM Insta, PhotoSearch; Force Motors: Predictive maintenance scale-up
3. Institutionalization	Enterprise-grade platforms, data governance, monitoring frameworks	Model explainability, change resistance, talent gaps	Cross-functional adoption, centralized data lake, MLOps setup	Tata AutoComp: Governance and retraining at scale; IndiaMART: Ethical AI frameworks
4. Business Transformation	Embedding AI across departments, measurable business outcomes, workforce upskilling	Security, compliance, drift management, unclear KPIs	Embedded AI in workflows, clear business KPIs, trust in AI outcomes	Force Motors: KPI-linked AI reviews; Maruti Suzuki: 26 digitized CX touchpoints
5. AI-Native Enterprise	Agentic AI, self-learning systems, real-time personalization, continuous innovation	Data quality at scale, GenAI risk, ROI clarity	Autonomous processes, adaptive learning, AI as a core capability	Verrock: Agentic AI readiness; Maruti: AI-powered customer assistants

"Customers don't want tools. They want intelligent solutions"

—Dr. Tapan Sahoo,
Head of Digital Enterprise, Information & Cyber Security, Maruti Suzuki India

functions. This gap between experimentation and enterprise-wide execution highlights a critical challenge.

As Indian enterprises increasingly adopt AI, the focus must shift from "How do we start?" to "How do we scale responsibly and sustainably?" The CIOs featured here prove that scaling AI is not a leap of faith—it's a disciplined climb. One that requires infrastructure, clean data, measurable goals, and cultural readiness. When done right, AI becomes not just a technology initiative but a business engine.

For India's CIOs, the opportunity lies in leading that climb—with clarity, control, and conviction. AI isn't the future of enterprise. It is the present, and the scale at which you deploy it will define the future of your business. Yet, as Deloitte's findings suggest, it will be a long slog, and realizing tangible ROI from AI will take more than a year—underscoring the need for long-term strategy, firm foundations, and operational integration. ■

Are You Building the Right AI Foundation—or Just Another Pilot?

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

Ashok Jade, CIO from Kirloskar Group in a conversation with CIO&Leader shares hard-won lessons on what it really takes to move AI from experimentation to enterprise-grade deployment.

MANY BUSINESSES are still having trouble scaling AI beyond pilots, even as it transitions from a buzzword to a boardroom priority. The CIO of Kirloskar Group, Ashok Jade, offers important insights on transforming AI prototypes into enterprise-ready solutions in this exclusive interview. He describes the practical difficulties faced by CIOs, ranging from handling model drift and fragmented data to negotiating old systems and vendor gaps. In a time when GenAI hype rules the roost, Jade's viewpoint serves as a reminder that effective AI is about governance, infrastructure, and long-term benefit rather than experimentation. Still, it's unclear if you're constructing the proper AI base or merely doing another little pilot.

CIO&Leader: What has been your most significant learning while moving AI projects from pilot to production?

ASHOK JADE: I feel, moving AI projects from pilot to production is a journey of recalibration of technical and strategic rigor.

One of the most significant learning is that model accuracy often deteriorates in real-world environments compared to controlled pilot settings. This drop is mainly because inconsistencies in data quality, misalignment between the model's assumptions and actual business workflows, and the evolving nature of user behaviour and operational contexts. The gaps show the gap between theoretical or limited thinking of problem and AI model with practical consumption.

Another critical learning is the underestimated cost of infrastructure. While pilots may run on shared or minimal resources, production-grade AI demands scalable, secure, and resilient infrastructure—often involving cloud orchestration, GPU acceleration, and real-time monitoring. These requirements can inflate operational costs significantly, challenging the projected ROI and necessitating a more careful while productionising AI projects.

One more learning is; pressing is the scarcity of skilled professionals and vendors with proven experience in large-scale AI deployments. The complexity of production environments calls for expertise not just in data science, but also in compliance, and business integration. Finding partners who can deliver optimized solutions at scale—without compromising on quality or cost—remains a formidable challenge.

These learnings underscore the need for a holistic AI strategy that goes beyond model development to encompass governance, scalability, and sustainable value creation..

Ultimately, de-risking AI isn't a one-time activity—it's a continuous discipline that evolves with the maturity of the organization and the complexity of the use case

CIO&Leader: What challenges did you encounter in scaling AI initiatives (e.g., data readiness, model drift, stakeholder buy-in, integration with legacy systems)?

ASHOK JADE: According to me scaling AI initiatives across an enterprise landscape is a complex combination of technical, organizational, and strategic challenges. One of the most persistent hurdles is data readiness—not just in terms of volume, but in quality, accessibility, and contextual relevance. Legacy systems often have data in silos, fragmented or unstructured, making integration with modern AI pipelines a non-trivial task. This fragmentation challenges the creation of unified datasets necessary for training AI models.

Model drift emerges as another critical concern, especially in dynamic environments where user behaviour, market conditions, or operational parameters evolve rapidly. frankly, stakeholder buy-in is often underestimated. While technical teams may be enthusiastic, business leaders and end-users require clear articulation of value, risk mitigation and alignment with strategic goals.

Integration with legacy systems further compounds the challenge. Many organizations operate on decades-old infrastructure that lacks the flexibility to support real-time AI inference or scalable deployment. Bridging this gap demands not only technical hurdle but also a phased modernization roadmap.

One more area of challenge is; governance, compliance must be embedded from the outset to ensure responsible scaling.

CIO&Leader: What best practices or frameworks have helped your de-risk and accelerate AI adoption?

ASHOK JADE: In our AI adoption journey, one of the most impactful learnings has been the importance of precise risk identification before any mitigation strategy is even considered. It's not enough to say "AI has risks"—we need to define which risks apply to which use cases. For example, when deploying AI on proprietary engineering drawings, the primary concern is intellectual property leakage. In contrast, when working with financial forecasting models, data security and regulatory compliance become the dominant risks. This contextual risk mapping has helped us avoid generic governance and instead build targeted safeguards.

Once risks are identified, we initiate structured cross-functional alignment. This involves bringing together stakeholders from IT, legal, compliance, and the relevant business function to validate the risks and agree on mitigation strategies. For instance, project involving predictive maintenance for industrial assets, one should align with operations and legal to ensure sensor data usage complied with internal data handling policies. Only after this alignment did you move forward with designing the solution.

The actual de-risking mecha-

In our organization, AI governance is not treated as an afterthought—it's embedded from the very beginning of each initiative. For every AI project, we establish an AI Steering Committee composed of both business and technology leaders.

nism is never just technical—it's a blend of technology and process.

Beyond these, we've also learned to de-risk through phased rollouts. Instead of deploying AI across the enterprise in one go, we start with a controlled environment—say, one plant or one business unit—where we can monitor performance, gather feedback, and refine the model. This not only reduces exposure but also builds internal confidence.

Vendor selection is another area where we've become more rigorous. We now prioritize partners who've demonstrated success in large-scale, cost-optimized AI implementations. We ask for references, proof of scalability, and even conduct joint pilots before committing to full-scale engagements.

We generally start small i.e., Pilot AI projects with clear ROI and manageable risk, we ensure we upskill our existing team who understand the AI and its risks, Invest in AI literacy across business and technical roles. We also use external audits to Validate models and processes through third-party reviews.

Ultimately, de-risking AI isn't a one-time activity—it's a continuous discipline that evolves with the maturity of the organization and the complexity of the use case.

CIO&Leader: How are you approaching governance, ethics, and monitoring of AI in production environments?

ASHOK JADE: In our organization, AI governance is not treated as an afterthought—it's embedded from the very beginning of each initiative. For every AI project, we establish an AI Steering Committee composed of both business and technology leaders. This committee is responsible for overseeing strategic alignment, risk posture, and ethical considerations throughout the lifecycle of the model. Ownership is clearly defined: each model has a designated owner who is accountable for its performance, compliance, and business impact. This clarity ensures that responsibility doesn't get diluted across teams.

From a compliance standpoint, our IT and security teams play a pivotal role. They maintain detailed documentation of training datasets, algorithms used, and the intended scope of each model. This not only supports internal transparency but also ensures we remain compliant with evolving regulations like GDPR and India's DPDPA. Before any model development begins, we conduct a formal risk identification exercise. For example, if the AI is being applied to proprietary engineering data, we assess IP risks; if it's used in finance, we focus on data privacy and regulatory exposure. A dedicated risk team tracks these assessments and escalates unresolved risks to senior management. ■

Should You Slow Down to Scale Right? CIO Answer to AI Hype

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

In an exclusive conversation between **Milind Khamkar, Group CIO at Super Max** with CIO&Leader, he reflects how intentionality beats intensity when deploying AI across complex, cost-sensitive operations.

MILIND KHAMKAR, Group CIO of a major consumer products company, provides a counterargument in an AI world that is characterized by excitement and speed: clarity should come before size. In this exclusive interview with CIO&Leader, Khamkar highlights the advantages of deliberate, ROI-focused adoption vs hasty experimentation. Given the limited constraints and intricate procedures, his team places a high priority on ethical AI, weighing use cases against scalability, cost, and business fit. His strategy strikes a mix between creativity and discipline, from piloting quality assurance and real-time OEE use cases to managing an enterprise-wide AI Ignite session. As interest in GenAI grows, Khamkar reminds us that preparedness, relevance, and governance are more important for success in AI than speed.

CIO&Leader: What has been your most significant learning while moving AI projects from pilot to production?

MILIND KHAMKAR: Our most significant learning has been the importance of discipline over speed. Given the current hype cycle surrounding AI, we've approached adoption with measured caution, recognizing it as a double-edged sword—full of potential, but also risk.

Being a low-margin, high-volume consumer product manufacturer, any technology introduced midstream must directly align with business outcomes. We've made it clear internally: only tangible success or measurable benefit justifies AI implementation. We're comfortable being late adopters, if it means the solution is sound, safe, and scalable.

We recently conducted a Global AI Ignite Session to raise awareness about responsible AI use. This initiative generated strong interest and user-proposed use cases, which are now being assessed for feasibility and value. Fortunately, we're not starting from scratch—we're building upon a strong foundation of automation and digitalization that has prepared us for this next leap.

CIO&Leader: What challenges did you encounter in scaling AI initiatives? (e.g., data readiness, model drift, stakeholder buy-in, integration with legacy systems)

MILIND KHAMKAR: We've faced the same set of challenges most businesses do—but layered with our own context:

- Our technology stack is deeply integrated, the result of years of evolution—not built for AI, but expected to support it.
- Data readiness isn't just about volume; it's about structuring, cleaning, and aligning data across diverse systems.
- Infrastructure needs triggered critical decisions around build vs. buy vs. hire—each with cost implications.
- Stakeholder alignment is a constant exercise. For a business driven by cost efficiency, every initiative must pass the Cost vs. Value test.
- Finally, identifying industry-relevant models—as opposed to generic AI tools—has been essential to avoid wasted effort.
- Scaling AI, therefore, is less about technology alone and more about enterprise fit, timing, and total cost of adoption.

CIO&Leader: What best practices or frameworks have helped you de-risk and accelerate AI adoption? Here's the approach we've followed to de-risk and build momentum:

MILIND KHAMKAR: Ignite Awareness: Conducted an enterprise-wide session on responsible AI usage to build foundational understanding and eliminate fear.

- **User-Led Exploration:** Encouraged business users to propose relevant use cases based on their operational pain points.
- **Start Small, Think Big:** Prioritized use cases with quick wins and potential for scalability.
- **Selective Tooling:** Evaluated and chose tools that align with our industry context and internal maturity. We're bypassing PoCs that repeat what the market has already proven.
- **Sustainable Setup:** Established clear sponsorship, funding models, and ownership, treating AI like a long-term capability, not a side project.
- **Governed Execution:** Created a feedback loop of learning, refining, and replicating success.

This framework ensures we balance experimentation with execution discipline.

CIO&Leader: Could you share a specific use case where AI delivered tangible business value for your organization?

MILIND KHAMKAR: We're in the early adoption phase, but several high-impact use cases are under consideration, including:

- **Quality Assurance Transformation:** Moving from sample-based checks to AI-powered full-product quality inspection.
- **OEE Monitoring in Real-Time:** Using AI to continuously monitor and improve overall equipment effectiveness (OEE) without manual intervention.
- We expect these initiatives to progress into production stages over the coming quarters with measurable ROI.

CIO&Leader: How are you approaching governance, ethics, and monitoring of AI in production environments?

MILIND KHAMKAR: We approach AI governance with the same seriousness as any critical IT initiative. Our guiding principles are:

Our AI governance is structured with clear roles and checkpoints to ensure accountability, while ethics is treated as a proactive discipline—not just compliance. Anticipating future regulations, we build transparently from the start. Our policies and training follow global best practices to ensure our AI systems are both effective and responsible..

A key mantra we follow: "The business existed without AI—our responsibility is to ensure it thrives safely with it."

As regulations evolve, we are committed to aligning with applicable frameworks and embedding AI accountability into our governance DNA. ■

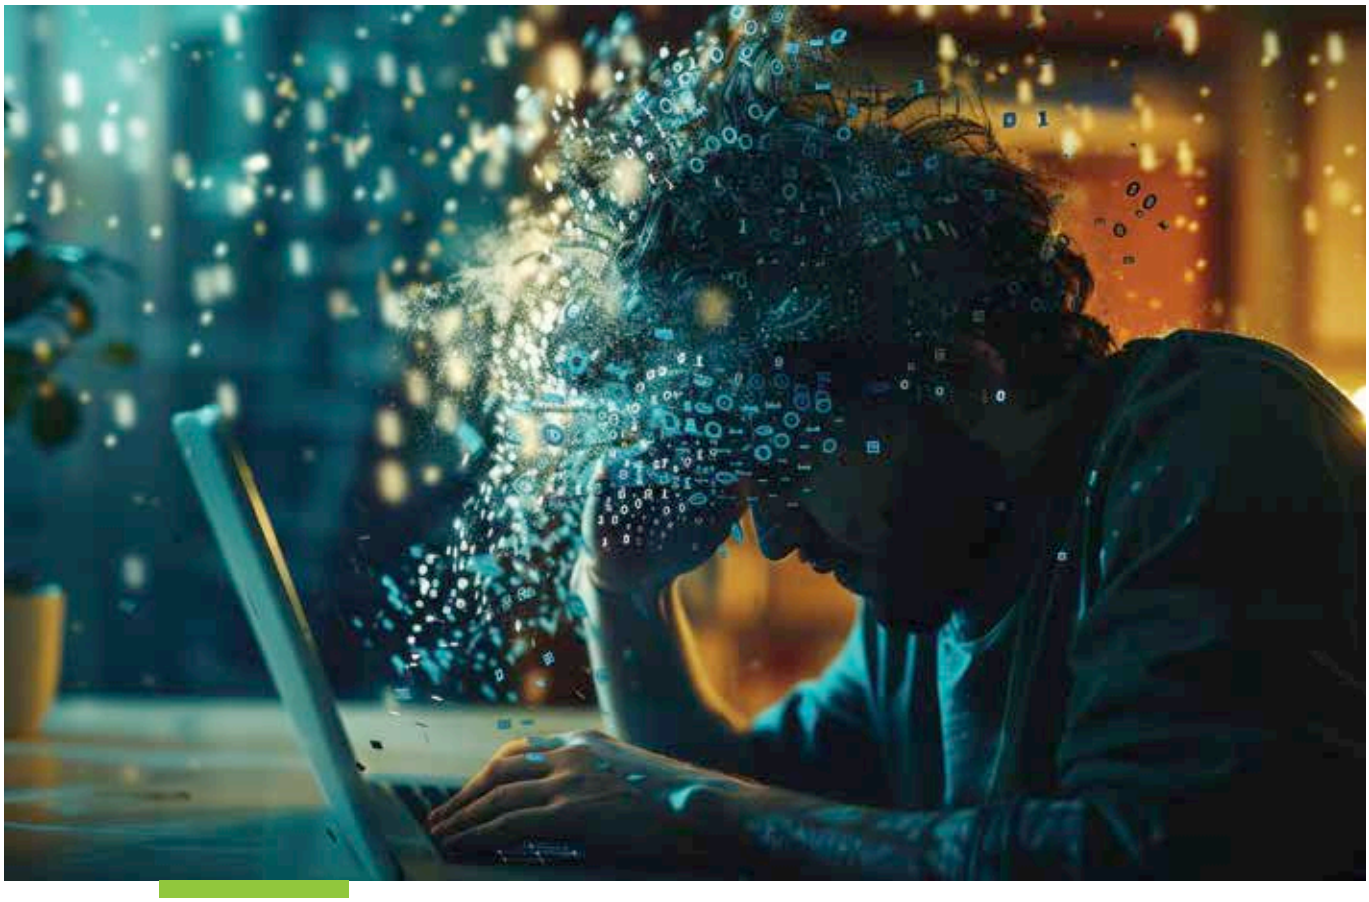

The digital battlefield: What's keeping security leaders awake in 2025

AI tops the list of cybersecurity fears in 2025, but outdated tools and missing basics still keep leaders up at night.

By **Jagrati Rakheja** | jagrati.rakheja@9dot9.in

IT'S 3 AM, and somewhere around the world, a cybersecurity professional gets an alert. Another potential breach. Another sleepless night. This scenario plays out thousands of times daily, as revealed in Arctic Wolf's comprehensive 2025 Trends Report, which surveyed 1,200 IT and security decision-makers across multiple countries.

AI Takes the Crown as Top Security Concern

For the first time in four years, artificial intelligence has overtaken ransomware as the #1 cybersecurity worry. Nearly 29% of security leaders now rank AI-related threats as their primary concern, pushing ransomware down to second place at 21%. This shift reflects growing

worries about AI-powered attacks that can craft more convincing phishing emails and identify system vulnerabilities more quickly than ever before.

However, experts warn against getting too distracted by AI hype. Traditional threats, such as malware, business email compromise, and ransomware, remain everyday realities that cause significant damage to organizations.

The Sobering Reality of Cyber Attacks

The numbers paint a concerning picture: 70% of organizations experienced at least one "significant cyber attack" in 2024. Even more alarming, only 25% of security leaders can confidently say their organization has not been breached – a sharp drop from 35% the previous year.

When attacks do succeed, the consequences are severe. Nearly two-thirds of significant cyberattacks resulted in productivity losses lasting at least three months, with some organizations experiencing disruptions for six months or longer.

The Ransomware Puzzle: Mixed Signals

While reported ransomware attacks decreased from 45% to 23% year-over-year, the threat remains serious for those affected. Surprisingly, 76% of ransomware victims still chose to pay the ransom, despite 90% working with professional negotiators who successfully reduced payments in 52% of cases.

What are the main reasons for paying? Preventing stolen data release (50%), speeding up recovery (49%), and having no alternative recovery method (48%).

Technology Gaps Persist Despite Heavy Investment

Despite 84% of organizations investing heavily in cybersecurity, significant challenges remain. A quarter of security leaders report outright dissatisfaction with their security tools, citing high false favorable rates (34%) and a lack of effectiveness (33%) as their top frustrations.

AI may be the new threat, but poor fundamentals remain the real enemy.

Surprisingly, AI security devices – despite the hype – deliver the least value, according to 18% of respondents, often generating more noise than useful alerts.

The Incident Response Revolution

One notable improvement: incident response (IR) preparedness has significantly enhanced. An impressive 88% of organizations now maintain IR retainers (up from 64% the previous year), and 81% have had to use them. This shows organizations are taking proactive steps to prepare for inevitable security incidents.

However, only 60% maintain current incident response plans, with just 35% having truly up-to-date documentation—a concerning gap in crisis preparedness.

Looking Ahead: Balancing Hype with Reality

As organizations navigate 2025's cybersecurity landscape, the key lesson is clear: while AI represents both opportunity and threat, fundamental security practices remain crucial. Multi-factor authentication, employee training, proper backups, and comprehensive incident response planning still form the backbone of effective cybersecurity.

The organizations that will thrive are those that embrace innovation while maintaining strong security fundamentals – because, in cybersecurity, yesterday's basics are today's lifelines. ■

The AI security tightrope: Business leaders race to innovate, security chiefs hit the brakes

As CEOs accelerate AI adoption, CISOs raise red flags—highlighting a growing divide between innovation ambition and cybersecurity readiness.

By **Jagrati Rakheja** | jagrati.rakheja@9dot9.in

IN CORPORATE boardrooms around the world, a fascinating drama is unfolding. On one side, CEOs are rushing to embrace artificial intelligence, viewing it as the golden key to future profits. On the other side, Chief Information Security Officers are waving red flags, warning that the AI revolution could become a cybersecurity nightmare.

New research from NTT DATA reveals just how wide this gap has become, with business leaders and security chiefs operating in almost parallel universes when it comes to AI adoption.

CEOs Betting Big on AI Revolution

Business leaders are placing massive bets on artificial intelligence, with 89% of CEOs identifying AI as the top technology needed to stay competitive. Nearly all organizations (99%) plan to continue investing in generative AI through 2026.

AI is revolutionizing industries from healthcare—where it analyzes medical data to predict patient outcomes—to manufacturing, where it streamlines production. Financial services use AI for fraud detection and intelligent customer service.

Security Chiefs Sound the Alarm

However, Chief Information Security Officers are putting the brakes on. While 94% of business leaders plan to increase security spending due to the adoption of AI, 88% of organizations express serious concerns about AI-related security risks.

The disconnect is stark: only 38% of CISOs agree their organization's AI and cybersecurity strategies are correctly aligned, compared to 51% of CEOs. Nearly half (45%) of security chiefs feel "pressured, threatened, or overwhelmed" by AI adoption—a sentiment shared by only 19% of other executives.

The Growing Threat Landscape

Security experts have good reason to be concerned. Cybercriminals are leveraging AI to craft sophisticated phishing emails and automate the discovery of software vulnerabilities. AI systems face unique threats, including adversarial attacks that manipulate input data, data poisoning that corrupts training information, and algorithmic bias that undermines system reliability.

We can't secure what we don't fully understand—and with AI, that gap is growing faster than ever.

Only 44% of executives believe their organizations can manage privacy risks from data poisoning attacks.

The Skills Gap Challenge

Adding to security leaders' concerns, 69% admit their teams lack the skills to work with rapidly evolving AI technology. Meanwhile, 72% of organizations still lack formal policies for AI usage, and 82% find government AI regulations unclear.

Finding the Balance

Security experts recommend enhancing AI visibility across organizations, developing comprehensive security policies, embedding security by design, prioritizing data protection, conducting rigorous testing, and maintaining continuous monitoring.

The challenge ahead is clear—organizations must harness AI's transformative power while building robust defenses against emerging risks. Success requires bridging the gap between innovation-focused executives and security-minded Chief Information Security Officers (CISOs). ■

Redesigning Enterprise Architecture for Scalable Intelligence

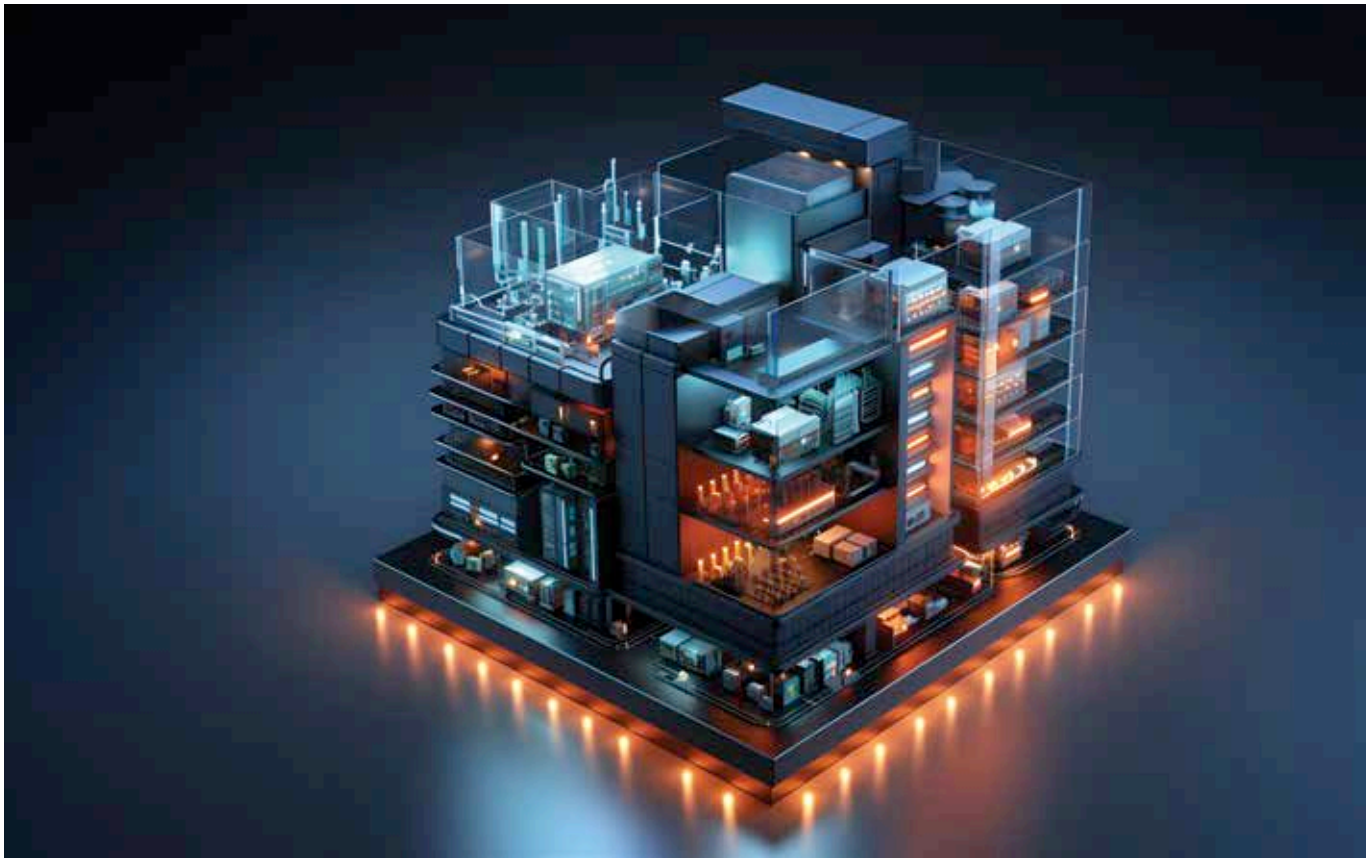

Agentic AI marks a new phase in enterprise automation, where context-aware AI agents act independently, adapt in real time, and operate across complex systems.

By **Niraj Kumar** | editor@cioandleader.com

ENTERPRISE AI is entering a new phase. The earlier wave of tools automated predictable tasks and assisted with decisions based on static models. But agentic AI is something far more dynamic. These context-aware AI agents can work independently, pursue goals, and adjust their behavior based on changing inputs. They are capable of initiating actions rather than simply responding to commands, which makes them useful in complex business environments.

Considering these advantages, over 80% of Indian organizations are now exploring the development of autonomous agents. However, deployment of these agents can be a challenge with traditional enterprise architecture, which was never built to support such autonomous systems. It means that the existing technology stack needs to be redesigned.

Building Scalable, Composable AI Pipelines

Earlier, scalability used to mean increasing server capacity or expanding cloud storage. But in the context of agentic AI, scalability means being able to manage operations across departments with the help of agentic AI. These autonomous agents need continuous access to structured and unstructured data. They also need clear rules on what they can and cannot do and a secure way to interact with both modern APIs and older systems still in use.

To make this work, organizations must move toward a composable model that facilitates the data-to-AI pipeline, automating the entire lifecycle from data ingestion to actionable AI insights. Information should also be structured in a way that allows agents to retrieve exactly what they need in real time without excessive computation or delays.

Embedding Continuous Security and Governance

In the past, human error used to be the primary point of failure. Now, with agentic AI, threats could come from faulty logic, compromised data, or external actors attempting to manipulate behavior. Therefore, modern enterprise architecture must adopt robust governance, observability, and security as core pillars supporting scalability.

This means designing systems with built-in traceability, rigorous access controls, and continuous real-time monitoring of KPIs and data drift. Furthermore, implementing zero-trust models ensures that every decision is verified in context. All actions should be logged, not just for compliance but for post-event learning. It is not just about stopping breaches but understanding what led to them.

Niraj Kumar
Group CTO
Onix

Ensuring Transparency and Explainability

Autonomous agents that operate in regulated sectors must not only work correctly but also be able to explain how they reached a conclusion. In that context, knowledge graphs and data lineage play a pivotal role. They enhance explainability and data visibility, providing clear, readable trails showing how decisions are made. For instance, if an agent rejects a loan application or reroutes inventory, the underlying reasoning should be transparent and accessible without requiring decoding of black box models.

Integrating Human Oversight and Control

AI agents are not here to replace human workers. They are here to handle the volume, repetition, and machine-speed reactions. People still define strategy, set goals, and manage relationships. For agentic AI to succeed, humans must remain part of the process. That means the architecture must include tools for humans to supervise, adjust, and occasionally override agentic decisions.

Enabling this kind of partnership requires thoughtful design. The interface must allow users to interpret, guide, and redirect agent behavior with ease. The training data itself must reflect the complexity of real-world scenarios to reduce unexpected outcomes. In that scenario, solutions like synthetic data generation can help build diverse, representative datasets that improve how agents perform under varying conditions. By combining machine efficiency with human judgment, organizations can unlock superior outcomes.

Building for What Comes Next

Agentic AI is not a passing phase. It is a structural shift in how digital systems work. As adoption increases, so will expectations. Enterprise architecture must evolve to handle not just smarter tools but smarter systems that learn, adapt, and act independently. Tasks such as code transformation and validation are increasingly automated, becoming integral to an evolving enterprise AI lifecycle.

This will require investment in addition to a cultural shift in how teams approach data, automation, and decision-making. But for those willing to make the change, the payoff is clear. Businesses will be better equipped to respond to change, scale operations, and unlock new value, driven by intelligence that grows alongside them. ■

Why Is Human Judgment Still Essential in AI Services?

AI amplifies intelligence by delivering speed and scale, but only human judgment can define purpose, navigate nuance, and ensure ethical, strategic decisions.

By **Joy Sharma** | editor@cioandleader.com

IN THE age of automation and algorithmic precision, it's easy to conflate speed with progress. Yet, the fundamental query remains as organizations are speeding up: Who charts the course?.

As artificial intelligence cannot determine purpose, it can accelerate decision-making.

In all our ventures, we view artificial intelligence not as a replacement but as an amplifier. It augments intelligence, scales output, and increases velocity. But judgment – real, contextual, human judgment – is what converts that velocity into vision.

AI Amplifies Intelligence, But Humans Define Purpose

Without a doubt, the AI systems of today are remarkable. They produce outputs at scale, identify patterns in milliseconds, and uncover insights more quickly than we ever could. AI can improve and speed up our progress, but it cannot determine our future. That's a tactic.

Strategy is not built on logic alone. It's built on intention. And intention requires context, ethics, long-term thinking, and a sense of impact – all of which originate in human judgment.

Furthermore, strategy is a reflection of values, intent, and long-term vision rather than something that can be created in a spreadsheet. The why will always be our question, even though AI can answer the what, how, and when.

Data Can't Make the Tough Calls

We often treat data like the ultimate truth – as if feeding enough of it into an AI system will magically give us the right answers. But it's not that simple. Data tells us what happened and some-

times what might happen. But it doesn't tell us why it matters or when it's time to challenge the logic.

In high-stakes environments – where regulation, culture, and public perception collide – nuance isn't optional; it's the whole game. AI can identify patterns and bring out insights, but human judgment is necessary to manage trade-offs, reinterpret rules, and take the tough calls.

Strategy seldom exists in absolutes. It exists in the gray, and however advanced our machines get, they're still learning to perceive beyond black and white.

AI Predicts Outcomes, Humans Understand Consequences

AI can forecast churn, model risk, or identify a market shift. But it doesn't grapple with the impact of those decisions. It doesn't inquire as to whether a behavior strengthens or weakens trust.

Human judgment is therefore essential. Long after the dashboards are cleared, we still bear the consequences of our choices in terms of our reputation, ethics, and the law. In business, speed is powerful. But foresight is what protects you.

Responsibility isn't a layer we add at the end. It's the foundation we build from the start.

Human Judgment Enables Adaptive Thinking

Though intelligent, AI has limitations. It operates based on what it has observed and learned. It doesn't rethink, imagine, or challenge the rules. Humans do. When circumstances change or the unexpected occurs, we adjust. We have the ability to question the plan, change course, and proceed with purpose rather than merely reason.

That's leadership, not fallback thinking. Seeing the wider picture rather than merely responding is the goal of adaptive thinking. And that necessitates reading more than just data; it involves reading the impact, the emotion, and the moment.

Joy Sharma
Founder & CEO, EZ

Humans Safeguard Ethical Boundaries in AI

One thing is clear: AI inherits the bias of its data. These systems reflect the world as it is—not as it should be. Without intention, they risk reinforcing the very gaps we aim to close. That's why human involvement is critical not just to validate outputs, but to question how they're built. We need people asking: What's the impact? Who gains, and who's left out? As algorithms shape decisions in finance, hiring, and healthcare, ethical thinking can't be occasional. It must be embedded in how we build, deploy, and lead every single day..

AI Accelerates Decisions, Humans Shape Outcomes

We talk a lot about decision velocity, how quickly we can move from insight to action. But speed alone isn't strategy. Great decisions demand context, foresight, and intent.

AI can analyze at scale, surface insights instantly, and flag emerging shifts. But it doesn't operate in isolation. Understanding why something matters, how a geopolitical event in one region might impact consumer behavior in another, or how timing and culture shape a product launch still requires a broader lens. Human leaders bring that lens. They add meaning to momentum.

Contextual Intelligence Is a Human Strength

While context is fluid, AI excels at pattern recognition. Why does a campaign thrive in one city but flop in another? Why doesn't success today guarantee loyalty tomorrow? These nuances stem from emotion, experience, and environment—not logic. True innovation often emerges from edge cases—those that defy patterns. And in such moments, human judgment matters most.

The Future Belongs to Human-AI Synergy, Not Supremacy

The goal of technology has always been to enable us to achieve more, not to replace what makes us human. It's about eliminating uncertainty, cutting down on inefficiencies, and making room for more in-depth, deliberate work. It's not a conflict between humans and machines; rather, it's about creating more intelligent systems that support and enhance one another.

For our global teams, AI is a trusted partner – one that extends our reach, sharpens our decisions, and moves with us at scale. And we know that the systems we're building today will define how businesses lead tomorrow.

The edge isn't artificial. The edge is how we bring human clarity and machine intelligence together with purpose, precision, and imagination.

Let's build on that. ■

How Far Can India's AI Engine Take It? A Look at Numbers

With $17 Billion on the Horizon, India's AI Economy is Gearing Up for Global Leadership

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

AI IS No Longer a Dream—It's India's Economic Reality. Artificial Intelligence is no longer an aspiration for the future—it's the driver of India's economic reality and digital future. Artificial Intelligence, once considered an arena for technology behemoths and deep R&D centers, now occupies the center stage of India's economic strategy, business processes, and public policy. In a 2025 Boston Consulting Group (BCG) report, "India's AI Leap:

BCG Perspective on Emerging Challengers," the domestic AI market in India is projected to triple its value—from levels at present to $17 billion by 2027. This is not prediction—it is warning. India is approaching a "leap" period in its AI maturity, during which talent, infrastructure, and enterprise aspiration are converging to build a competitive AI economy on a global scale.

Talent Is India's Competitive Advantage

India's AI ascension is rooted in its large and expanding talent pool. With more than 600,000 AI professionals currently working, and estimates to double up to 1.25 million by 2027, India dominates close to 16% of global AI talent base—second only to the United States. This talent pool provides quality, affordable expertise that enables Indian companies to develop and implement enterprise-class AI at scale and speed. For international businesses, it places India not only as a delivery hub, but a center of AI innovation.

Infrastructure, Startups, and Public Platforms Power Growth

India's public digital infrastructure is facilitating real-time, scalable AI deployment across verticals. Platforms such as Aadhaar, UPI, and ONDC offer interoperable data rails, while 700+ million internet users provide unparalleled scale for digital experimentation. The emergence of more than 2,000 AI startups in a span of three years is a testament to the entrepreneurial depth of India's innovation ecosystem. Government encouragement through the ₹10,000 crore IndiaAI Mission and collaborations with international leaders such as Intel are increasing access to compute, skills, and strategy.

India possesses the talent, the tools, and the momentum. What India needs now is visionary, bold leadership to shift this momentum into long-term gain.

Enterprise AI in Action: Pilots to Production

Indian companies are not sitting on the sidelines—they are infusing AI into core business operations and redefining business models:

Razorpay leverages AI to automate almost 80% of customer onboarding, enhancing speed and accuracy. Lenskart employs facial recognition and mapping to build a personalized virtual try-on for eyewear. Pocket FM used generative AI to reduce content creation costs by 90% and created more than 100,000 hours of audio content. NoBroker utilized AI-driven solutions such as Iris for image authentication, speeding up property listing operations.

These examples show that AI can scale from experimentation to enterprise-level ROI.

BCG's Roadmap: How to Scale AI Responsibly

As per BCG, successful AI implementation depends not on tools but on organizational design. Its 70/20/10 model recommends that companies invest in:

- 70% in people and processes
- 20% in data infrastructure
- 10% in tools and models

This is enabled by five essential enablers:

- Cloud-native, modern tech stacks
- Cross-functional teams with combined business and tech
- AI Centers of Excellence (CoEs)
- Workforce reskilling and behavior nudges
- Ethical AI governance and transparency mechanisms

Barriers to Break: Talent Shortfalls and Legacy Systems

Though growth is bright, the journey is not smooth. A potential talent shortfall of more than 1 million AI professionals by 2027 threatens to haunt companies, says Bain & Company. Most Indian businesses continue to run legacy IT stacks that weren't built to handle autonomous, context-aware AI agents. Additionally, concerns around data bias, governance gaps, and unexplainability need to be overcome as AI systems scale.

- C-Suite Action Plan: From Pilots to Purpose
- Business leaders, it is time to act. In order to compete on the world stage and lead locally, C-suites need to:
- Re-architect systems for AI-readiness
- Develop internal governance teams and ethics frameworks
- Invest in ongoing reskilling and experimentation
- Prioritize value realization over feature adoption

Leadership needs to shift from tooling up to owning smart decision-making at scale.

India's AI Future Is Now

According to Mandeep Kohli, Managing Director and Partner, BCG India:

"AI is no longer optional—it's foundational."

India possesses the talent, the tools, and the momentum. What India needs now is visionary, bold leadership to shift this momentum into long-term gain. With the right combination of innovation and governance, India can set the gold standard for AI economies globally. ■

Ashwani Narang
Head of Finance & Spend Management—India Subcontinent
SAP

How SAP is Localizing Innovation and Leading Intelligent Enterprise Shift

Ashwani Narang, Head of Finance & Spend Management—India Subcontinent, SAP discusses how SAP is localizing AI, cloud, and compliance to accelerate India's intelligent enterprise journey.

By **Jagrati Rakheja** | By jagrati.rakheja@9dot9.in

AS INDIAN enterprises focus on AI, cloud-native models, and sustainable growth, SAP aligns its strategy with the nation's digital ambitions. In this exclusive interview, Ashwani Narang, Head of Finance & Spend Management—India Subcontinent at SAP, discusses how the company enables enterprises to innovate faster, operate more efficiently, and remain compliant with evolving regulations.

From embedding generative and agentic AI into core applications like finance, HR, and supply chain to building API-driven integrations with initiatives like ONDC and GST, SAP drives meaningful outcomes at scale. Narang also highlights the significance of SAP's ₹500 crore India Data Center investment and how nearly 25% of the company's new patents now originate from India. With a new R&D campus underway, SAP is strengthening its commitment to innovation, localization, and customer-centric transformation, empowering Indian enterprises to become truly intelligent and future-ready.

CIO&Leader: What role do you play in ensuring a strong AI-driven foundation, particularly in the context of SAP India's data center infrastructure?

ASHWANI NARANG: The roadmap for us is clear on AI. Now, if you look at it—if I draw a big sphere for you—you would see that AI is right at the top. You have ML, which is machine learning, as a subset; then you have deep learning, which is more focused on neural networks; and finally, we are getting down to generative AI. The roadmap is pretty straightforward for us.

We are looking at combining generative AI and agentic AI. When we consider agentic AI, we see use cases, and if you were upstairs today, you would have heard about 180+ use cases that already exist. We are speaking to customers; some customers have already gone live with it. We see specific use cases we have already built when considering generative AI. Even SAP has developed some generative use cases internally.

For example, if you're a new employee in your company and you want to set up goals for that particular employee, then essentially what we are talking about is: I, as a journalist, for example, enter

the organization, and the system should predict what my KPIs should be. Now, that would be awesome—if the system could say, "These are the top five KPIs for journalists working in technology." That's generative AI for you.

Now, when you talk about that in the context of India's data center, we're saying that SAP India has already pledged a certain amount of money. It was reported in the media some years back—about ₹500 crore—to build the local data infrastructure. The India data center is already live for the entire ERP ecosystem. Regarding lines of business, it is live for Ariba; regarding hire-to-retire processes, it is live for SuccessFactors. So, we are integrating all these capabilities from both an application and an AI standpoint. We're building that into the India data center today. The roadmap is clear—what happens globally will also occur in India. SAP has pledged that the entire set of use cases will be implemented in the India data center.

And one more thing I would like to add—you may or may not have heard it—is that about 25% of SAP patents on new technology are now being generated from India. That is the highest percentage outside our headquarters in Germany. A new R&D campus is coming up, including new labs and facilities, and we plan to employ more people there as well. That's how committed we are to AI in India.

CIO&Leader: How are you enabling your customers to accelerate their digital journeys by harnessing the full potential of the cloud to innovate faster, operate more efficiently, and achieve their long-term growth strategic goals?

ASHWANI NARANG: If you remember the good old world—the good old world talked about ERP on-prem—and those days are gone now. Cloud will bring you a new set of innovations very clearly, which is pretty evident. The innovation happening—not only on our side but in any technology company—is happening on the cloud. As a customer, when you empower yourself with the SAP cloud business, you will benefit from all the innovations being integrated into all these areas.

We are building the AI layer on top of our enterprise applications—AI will make life easier for companies.

For example, if you are a procurement manager in a company and want to access the system and see how many pending POs you need to finish as your to-do task, I have an agentic AI on top. Can you show me a list of the pending POs? And here it goes—you have the entire list available in five seconds.

Or if you're a casual user and have no idea about the policies of the company you're working for—because it can be a significant document, by the way—we can build that entire document into the system. The system can then tell you: "Look, you are a manager, you've joined this company, and you cannot buy a laptop worth more than ₹50,000 because that's the policy set up by your organization." So, in the system, while you see three laptops, only two are eligible for you to buy because they are priced below ₹50,000. The system will stop you then and there if you try to violate the policy, knowingly or unknowingly.

Therefore, all of this intelligence is being built into the system. Customers are using this now to ensure compliance and reliability and to work within their organization's policy framework. That is what makes the difference for all these companies.

CIO&Leader: How is AI fueling SAP growth in India? What are the key cost and complexity challenges it addresses for enterprises?

ASHWANI NARANG: If 10 tech startups are happening around the globe, almost 7 are native AI, which also holds for India. So, as a provider, the obligation is on us, as SAP, to go back to our customers and say: we will give you the best-in-class AI experience. And that's an obligation we have taken on ourselves.

The crux of the matter is how we are helping our customers transform better. As I said, we already provide our customers with about 180 AI use cases. There was one, you know, from a public sector company that I remember, which asked us: "If I want to go out and close my books, is there a way that Joule can help me do that?" We saw that upstairs at SAP Now today.

But then, what I'm trying to get to is that AI will make life easier for companies. So, whether it's the new-age companies or companies like SAP, we are building the AI layer on top of our enterprise applications. On the other hand, the new-age companies are also starting from that point. AI will be evident, and it's now on us—how fast-paced we can be to ensure that we give the best-in-class experience back to our customers.

CIO&Leader: How is your traditional on-premise model evolving to a cloud-first approach, and what does that transition mean for your customers in India?

ASHWANI NARANG: I remember

The customer value journey starts when we justify the business case—once ROI is clear, the cloud shift becomes obvious.

my conversation with one of my customers almost one and a half decades ago, 15 years ago. I asked them, "You know SAP has launched S/4HANA, do you know this?" And they said, "Yeah, yeah, we know it." So I asked, "When are you converting from ECC to S/4HANA?" And they said, "You know what, we are happy with what we are doing now. Probably it's not the time for us to move from ECC to S/4HANA."

The single biggest problem was that we were living in the mindset of "I am already doing a fair job—why should I change?" And that doesn't hold good today. If you look at the evolution of technology or the disruption happening globally right now, if you do not have the best-in-class technology, you're going to get stuck.

For example, if I'm a pharma company and I have a shipment coming in—from China to India, Germany to India, or even the US to India—I need to track a path that ensures the shipment arrives on time. That's a fair expectation. But if I want to know where my shipment is today—whether it is at sea, which particular sea, and at what location—then think back to COVID.

A drug coming in from one place to another had to be maintained at a specific temperature and pressure. Therefore, I need to know if the conditions are being met and if the route is followed correctly. If the temperature and pressure are not maintained, the drug will not be effective, and that's a loss of money, and possibly lives as well.

That's why SAP realized we need technology that operates in real time. We need to help our customers be more predictable in how they run their businesses, and businesses understand that.

So, while we returned to customers, they returned to us and said, "We understand on-prem was a thing of the past. Help us understand what advantages we will get if we move from on-prem to cloud. What's the impact? Will you help us reduce costs? Will you help us become more effective and efficient in completing our work faster?"

We prepared all those business cases for each customer. Customers found it easier to decide once the business value was proven to them: "Well, this makes sense. I'm getting my ROI." And so, it became the right time to move from on-prem to cloud.

That is how we are creating a journey. Internally, we have a name for it—the customer value journey. That's what we call it. The customer value journey starts when we justify the business case to them. I think the ROI convinced customers that there is value in moving from on-prem to cloud. Of course, we are going to customers, but customers are also returning to us. So, it's a win-win for all of us.

CIO&Leader: As Indian enterprises increasingly adopt AI and cloud-native models, how is SAP collaborating with the local ecosystem—startups, partners, and developers—to drive innovation at scale?

ASHWANI NARANG: I love this question because there must be something in it for the local MSMEs and the economy.

Now, when we talk about how SAP is collaborating—or creating those APIs where two systems can connect—these systems may have something to do with a solution SAP has built, or perhaps it's an extension we've built with a partner, a customer, or even in collaboration with a government marketplace.

I'm thrilled to give an example. The other day, I was sitting with the Chairman of ONDC. You know, ONDC is a significant initiative by the Government of India, and they

want to bring every small MSME onto a platform where they can access a marketplace. This marketplace connects with SAP, enabling buyers and sellers to trade. The question posed to us was: "SAP, can you connect to ONDC? And if you can, can we have a marketplace where all the MSME providers can participate and help evolve the Indian economy?"

I'm proud to tell you that we took that challenge upon ourselves and are actively working with a few government-sector companies to identify how we can integrate with ONDC.

Suppose someone like me in a PSU wants to get on a system and buy something. In that case, I can connect to ONDC and make that purchase, where the supplier doesn't pay anything to the platform because the Government of India supports it in a charitable format to help MSMEs grow.

We've also developed extensions with local partners. For example, let's discuss GST, a significant concern for CFOs. They often worry: If I've paid GST to a supplier, has that supplier submitted it to the government on time? The law says the buyer becomes liable if the supplier hasn't done so. So, I, as a customer, am concerned that everyone complies with the law and submits taxes on time.

We've developed a specific API with one of our partners in India that can fetch data from the government's open-source portal and alert me. It can tell me: "A supplier is collecting GST from you on time, but they're submitting it to the government late—40% of the time." That delay makes me liable. This is another innovation we've taken to the market, and customers have very well received it.

There are more examples I can talk about, but these two were particularly relevant.

CIO&Leader: With compliance, data sovereignty, and sustainability gaining importance, how is SAP aligning its India strategy to meet these evolving enterprise priorities through its infrastructure and services?

ASHWANI NARANG: So, compliance and sustainability—I was looking at a CEO survey conducted by Gartner, and compliance was listed as point three, acknowledged by 40% of CEOs as essential. Sustainability ranked at point seven, cited by almost 25% of CEOs who said they need sustainable practices.

Again, SAP hits its obligation to develop compliant systems, protect our customers' interests, and be forward-looking, ensuring we also address the sustainability agenda.

When I meet a CFO, I ask them: You prepare your profit and loss statement every quarter, report it to the street, SEBI demands it, you're a public company, etc. But you're at the same time, do you have your green ledger? While you may be doing well regarding your profit and loss, how are you doing regarding sustainability?

For example, if I ask you a simple question—if you're spending ₹100 in your company, what percentage of that is going to a sustainable supplier? Can you confidently say that 10% of your spending is toward sustainable suppliers? These are the kinds of questions that are now being asked.

Even from a regulatory perspective, BRSR (Business Responsibility and Sustainability Reporting) must now be filed, just like you declare your P&L results every quarter, you must also submit your BRSR results quarterly. If you don't, you could be excluded from mid-don'tndices, your valuation may drop, and your stock price could plummet. That's not a good sign.

So, if I'm a toThat's any to's any, I need to ensure my I'm SR score is in place. I need to generate and submit that report to the government, just like I do with my P&L. That has become critical.

Companies regularly ask us: How can you help us make our processes more sustainable while ensuring we comply with the regulatory frameworks defined within our organizations?

We've already developed and backed your Scope 3 emissions there. We have a product called SDX—Supplier Data Exchange.

For example, if I issue a PO to a supplier and ask them to deliver 10 items, they may say "yes," but then I'll ask, "Can you tell me our Scope? I'll know about each line item of these 10 products?" And that's the kind of challenge we're facing.

For example, I live in Delhi and flew to Mumbai today. How many carbon emissions did the flight generate? That's a great question.

According to the UN framework on sustainability, companies must fulfil about 19 criteria. In the function I handle—finance and spend management—we address 13 of them. One of the criteria is to measure how much carbon is generated by employees in your organization.

As we say, charity starts at home, so the change has to start internally. I need to be sustainable myself before I can advocate for company-wide sustainability.

We've provided customers with tools and innovations to say, "Don't worry—we're looking at your greenhouse gas emissions, tracking employee-generated carbon emissions, and measuring Scope 3 emissions." We're a line of products that can help you become compliant and sustainable wherever you are. That is how we are approaching it.

CIO&Leader: What are India's top priorities, including revenue goals

from Indian operations and innovation focus areas?

ASHWANI NARANG: Okay, so revenue, of course—we look at it globally. The numbers are already explained at the global level. The results were out yesterday, and we are delighted and very gung-ho. Everything is looking green, so I'm pretty excited about it. The share price also tells a story about it.

But in terms of innovation, there is a lot of localization effort going on—that's where I would love to bring your attention. When you look at the kind of use cases we're getting from a localized standpoint—for example, if I'm invoicing with a particular supplier—there is the IRN portal of the Government of India. The invoice must be generated in five copies, barcoded, and returned to my system.

This means that the system should be capable of handling these invoices and retrieving them from the IRN portal after submission. And that's what we've done. Few providers in the country are capable of doing that. That's one example of localization or the kind of local innovation we are bringing to the table.

Another example is that Indian organizations often use a simple process called Note for Approval. Many companies follow this process while signing off on a capital expenditure, and follow a particular format and process. We are customizing all of this to meet customer needs on our BTP platform.

Now there's another aspect—scaling it to an unimaginable level. Customers looking for customizations or configurations beyond SAP's core product offering, inside our portfolio, are now getting extensions developed on BTP. So I think, in many ways, there are a lot of innovations planned from a localization standpoint, and even very specific to individual customer use cases.

We don't just call it AI—we call it Business AI. It must be reliable, relevant, and rooted in real use cases.

CIO&Leader: With SAP integrating AI into finance, supply chain, and HR, how is the platform helping CIOs balance agentic AI autonomy with human accountability in critical enterprise decisions?

ASHWANI NARANG: We are looking at CFOs, CIOs, CPOs, Chief Supply Chain Officers, and Chief HR Officers. You know, all of them have their own set of requirements. When we're building AI—or specifically generative AI—use cases, we consider many underlying factors. For example, you touched upon security.

Our philosophy is quite clear—we call it Business AI. We're not just calling it AI—we call it Business AI. It has to be reliable, relevant, and related to the specific use case we're working on. That's our definition of Business AI.

We understand these personas very well when looking at those kinds of datasets.

Regarding security and data protection, we've already obtained the necessary SOC approvals and certifications, and we are following all those protocols to ensure customer data is safe and the AI is secure. We're aligned with global guidelines on artificial intelligence.

So, I think that becomes very important—and it is at the core, the central theme of SAP's approach to building all its AI products.

We're doing this by bringing to life many customer use cases across different personas and CXOs. ■

Palanivel Saravanan
Vice President of Cloud Engineering at Oracle India

Building smarter AI begins with better data

Palanivel Saravanan, Vice President of Cloud Engineering at Oracle India, on AI, cloud expansion plans, and market differentiation

By **Jatinder Singh** | jatinder.singh@9dot9.in

AS ONE of the world's leading cloud infrastructure and enterprise software providers, Oracle has played a pivotal role in accelerating digital transformation across industries. In India—a market experiencing exponential digital growth—Oracle is strategically expanding its cloud footprint to address rising demands around data sovereignty, regulatory compliance, and AI-driven innovation.

With dual cloud regions in Mumbai and Hyderabad and deep investments in AI-native infrastructure, Oracle is positioning itself not just as a global cloud provider but one deeply integrated into the Indian ecosystem.

In this exclusive interaction, Palanivel Saravanan, Vice President of Cloud Engineering at Oracle India, shares insights into Oracle's expansion in India, its differentiation in a hyperscaler-dominated market, and how it's enabling enterprises with AI-native solutions, industry-specific services, and future-ready cloud engineering capabilities.

CIO&Leader: Oracle has steadily expanded its cloud footprint in India. What are the key drivers behind this growth?

PALANIVEL SARAVANAN: India is one of the fastest-growing digital economies and a key market for Oracle. Our cloud expansion is driven by the evolving needs of enterprises, especially around data sovereignty, regulatory compliance, and latency-sensitive workloads. These are non-negotiable, particularly for public sector and regulated industries.

Our dual-region cloud strategy in Mumbai and Hyderabad ensures organizations can store and process data locally, meeting stringent compliance and data residency norms. Additionally, Oracle's AI-native infrastructure, including a supercluster of 130,000 GPUs—one of the world's largest—caters to Indian enterprises' performance and compliance needs without relying on overseas data centers.

We embed AI capabilities across all layers—databases, platforms, and SaaS applications—reducing the need for separate AI infrastructure. Innovations like Select AI, MySQL HeatWave GenAI, and Oracle Database 23ai deliver real-time insights and intelligent automation, enabling agility, security, and scalability.

CIO&Leader: How does Oracle differentiate itself from hyperscalers like AWS, Azure, and Google Cloud?

PALANIVEL SARAVANAN: Oracle's competitive edge lies in its AI-centric design philosophy and deep integration across the stack. Unlike vendors offering standalone AI services, we embed AI within its infrastructure, databases, and SaaS platforms, enabling faster value realization.

Our AI-ready databases like MySQL HeatWave GenAI and Oracle Database 23ai bring LLMs and vector processing to where the data resides—eliminating the need to move data, which boosts performance and enhances security. This is particularly effective in hybrid and multicloud environments.

With Select AI, users can query databases in over 50 languages using natural language, democratizing access to insights. We also offer prebuilt AI agents within Fusion Applications for functions like customer service and compliance, as well as support for customer-owned LLMs—giving enterprises control across multicloud setups.

CIO&Leader: How are you helping CIOs manage cost transparency, avoid vendor lock-in, and ensure architectural flexibility?

PALANIVEL SARAVANAN: We address CIO concerns with an AI-native, open architecture that emphasizes portability, cost efficiency, and autonomy. Oracle Database 23ai introduces autonomous capabilities like self-tuning and self-patching, reducing manual oversight and long-term costs. Our platforms support open standards and allow easy deployment of enterprise workloads and AI models using tools like AI Quick Actions. This flexibility empowers businesses to customize AI without dependency on external infrastructure or niche expertise.

In hybrid and multicloud environments, MySQL HeatWave GenAI brings LLMs into the database layer, minimizing egress costs and enabling real-time analysis of structured and unstructured data. Oracle also provides vector search for advanced analytics, especially useful in customer feedback, IoT, and personalization.

We embed AI capabilities across all layers—databases, platforms, and SaaS applications—reducing the need for separate AI infrastructure.

CIO&Leader: How are you integrating generative AI and machine learning to enable digital transformation?

PALANIVEL SARAVANAN: Our guiding principle—"AI works best when the data is best"—informs our strategy. Oracle doesn't treat AI as an add-on. We've built it into our infrastructure, platforms, and applications for years, well before the GenAI boom.

At the heart of this is Oracle Autonomous Database and Oracle Database 23ai, which include over 300 AI features automating tasks such as patching, performance scaling, and workload optimization.

Select AI allows users to query data in natural language—eliminating the need for SQL. HeatWave GenAI integrates LLMs and vector search directly into MySQL, providing deep insights without moving data. This ensures compliance, agility, and reduced complexity.

AI is also embedded in our SaaS applications—ERP, HCM, SCM—supporting use cases like predictive maintenance, safety compliance, and dynamic customer insights.

CIO&Leader: Which sectors are seeing the highest adoption, and what use cases stand out?

PALANIVEL SARAVANAN: We see strong traction across financial services, retail, education, and public sector. Indian enterprises demand domain-specific intelligence, and Oracle delivers through tailored solutions.

For instance, Muthoot FinCorp migrated its on-prem ERP system to OCI, improving system performance and operational efficiency by 50%. In banking, Select AI helps teams access customer insights instantly, simplifying service delivery and compliance.

In the public sector, Oracle supports RAG (Retrieval-Augmented Generation), enabling contextual AI insights using proprietary datasets. The Ministry of Education chose OCI to modernize DIKSHA, India's national ed-tech platform. The move improved accessibility, reduced IT costs, and now supports 1.48 million schools in 36 languages.

Across sectors, Oracle's AI-native cloud helps businesses reduce costs, improve decision-making, and enhance citizen or customer experiences.

CIO&Leader: What's next for cloud engineering, and how should CIOs prepare?

PALANIVEL SARAVANAN: As AI becomes foundational, cloud engineering is evolving into a strategic function. Tomorrow's cloud teams must move beyond infrastructure to master AI, automation, and integration.

CIOs should focus on developing AI-native capabilities—from LLM development and vector database management to secure RAG workflows. Oracle supports this through AI Quick Actions, embedded agents, and custom LLM support.

Equally critical is a platform-first mindset. Engineers must integrate compute, storage, AI, and SaaS within hybrid and multicloud environments.

We're actively supporting this shift with training programs, certifications, and CoEs across India. Future cloud teams must be as fluent in AI and automation as they are in provisioning infrastructure—turning IT into a business catalyst. ■

Where CIOs Lead, Conversations Thrive

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

Total engagement metrics

Last 90 days

1	CIOandLeader (Your Page)	47,730	▲69.5%
2	ETCIO	26,972	▲53.8%
3	ETCISO	16,269	▲30.4%
4	Information Security Media Group (ISMG)	15,485	▲233.6%
5	CISO FORUM	6,125	▲118%
6	Express Computer	5,833	▲11.8%
7	Dataquest	4,785	▲9.9%
8	ETTelecom	983	▼84.4%

Source: LinkedIn Analytics I Time range: 1 April 2025 - 29 June 2025.
Total engagement metrics: Reactions, comments & shares.
Percentage shows change from the previous 90 days.

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

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

RNI No. DELENG/2012/42736

AMD PRESENTS

26th Annual Conference CIO&LEADER

AI: From Pilot to Production

1st-3rd August, 2025 • Ritz Carlton, Pune

CO-PRESENTED BY NxtGen

#CIOandLeaderConference

The 26th Annual CIO&Leader Conference is not just another tech gathering. It is a curated leadership experience where India's most forward-thinking CIOs converge to discuss, debate, and drive the transition of AI from experimentation to enterprise-wide impact. Over three dynamic days, the conference will bring together 200+ top CIOs and IT leaders, global AI experts, and many more.

Eminent Speakers @26th Annual CIO&Leader Conference

Guest of Honour

Shri Adv Ashish Shelar

Minister for Information Technology & Cultural Affairs, Government of Maharashtra

Partha Iyengar
Ex-Gartner Fellow, Country Manager Research (India), Gartner, HBR Author, Ex Board Member SBI, Speaker, Advisor

Prof Balaraman Ravindran
Head: Wadhwani School of Data Science & AI / Robert Bosch Centre for Data Science & AI / Centre for Responsible AI, IIT Madras

Anil Kumble
PadmaShri awardee and former captain & coach of the Indian cricket team

Annu Kapoor
Actor, singer, director, radio disc jockey and television presenter

Dr. Nitish Bharadwaj
Director, screenplay writer, actor, orator, and spiritual thinker

Chaitanya Devadhe
Indian Idol Season 15 finalist

Mohan Radhakrishnan
Acknowledged to be a voice reminiscent of Kishore Kumar

Know more: **https://annualconference.cioandleader.com/** **Follow us:** @CIOandLeader

For partnership inquiries, write to **Hafeez Shaikh, hafeez.shaikh@9dot9.in, +91 9833103611**

PRESENTING PARTNER: AMD | CO-PRESENTING PARTNER: NxtGen | CO-POWERED BY PARTNERS: freshworks, TATA COMMUNICATIONS, VERTIV. | GOLD PARTNERS: SentinelOne, UNICOM, neat.

ASSOCIATE PARTNERS: DELTA, neysa | EXHIBIT PARTNERS: Barracuda, infraon | MEDIA PARTNER: CNBC TV18 | CONCEPT BY: CIO&LEADER, IT NEXT, 9100