APRIL 2026
SUPPLEMENT OF CIO&LEADER

A 9.9 GROUP PUBLICATION
www.cisoforum.in
cisoforum-in

CISOFORUM

Security For Growth And Governance

WHEN AGENTS GO ROGUE

Autonomous AI is already making decisions that no one approved, and no one can fully trace

Abhishek Jha
Tata Technologies

Dr Pawan Chawla
Tata AIA Life Insurance

Himachal Jothinarasimhan
Ashok Leyland

SPECIAL

Outlook 2026

How Cybersecurity Leaders in India Are Navigating Digital Transformation, Regulatory Complexity, and Emerging Threats

CIO&LEADER studiotalks

CIO&LEADER STUDIOTALKS—
WHERE TECHNOLOGY MEETS THE SPOTLIGHT!

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

WHY JOIN STUDIOTALKS?

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

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

Be featured among India's top technology leaders.

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

SECURE YOUR SPOT NOW!

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

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

Follow us: @CIOandLeader

A New Reckoning

THERE IS a particular kind of silence that falls over an industry when it realizes the ground rules have changed. Security professionals who have spent years deploying layered defenses, negotiating patch cycles, and educating boards found themselves reaching for language that didn't quite exist yet — not because the threat was incomprehensible, but because it was, finally, undeniable. The inflection points which researchers had been quietly anticipating had arrived, and it had a name - Mythos.

Mythos is not merely a more powerful tool. It is the moment when offensive capability becomes ambient — when the expertise required to conduct a sophisticated, multi-stage cyberattack stops being a scarce human resource and becomes a prompt. This shift demands a response, and it needs one now.

Yet even as we absorb the implications of AI-driven exploitation, a second storm is gathering — one that moves more slowly but is potentially more destructive. Quantum computing is here, and the timeline to cryptographically relevant systems has compressed from a theoretical "sometime this century" to the next decade. RSA, ECC, and the public-key infrastructure securing virtually every enterprise transaction will be rendered obsolete when a sufficiently powerful quantum computer is pointed at them. Nation-state actors are already harvesting encrypted data today, intending to decrypt it tomorrow with quantum systems.

For CISOs, the convergence of these two forces creates a threat environment of unusual complexity that requires a two-speed response. The AI crisis is an immediate operational urgency: compress patch windows, rebuild incident response around hours, and equip teams with AI-augmented defenses. The quantum crisis requires strategic patience: audit your cryptographic estate, identify long-lived sensitive data, and begin migration to post-quantum standards.

The emerging imperative is not technical alone. It is organizational, financial, and cultural. Boards must understand that past risk models are already obsolete. We are living through a compression of threat timelines with no modern precedent. AI is collapsing the present.

Quantum will be quietly dismantling in the future. The organizations that will navigate this era are not those with the largest budgets, but those with the discipline to act before the breach rather than in its aftermath. ■

"The organizations that will navigate this era are not those with the largest budgets, but those with the discipline to act before the breach rather than in its aftermath."

R. Giridhar
Editorial Director,
Enterprise Tech Publications
c-raja.giridhar@timesgroup.com

CONTENTS APRIL 2026

COVER STORY

Cover Design by:
Manish Kumar

Please Recycle This Magazine And Remove Inserts Before Recycling

NEWS & VIEWS

SURVEY

INSIGHTS

OPINION

CISOFORUM
Security For Growth And Governance
www.cisoforum.in

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

EDITORIAL
Group Editor: **R Giridhar**
Editor: **Jatinder Singh**
Senior Correspondent & Editorial Coordinator – CISO Forum: **Jagrati Rakheja**

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

SALES & MARKETING
Senior Director - B2B Tech: **Vandana Chauhan**
Head - Brand & Strategy: **Rajiv Pathak**

National Sales Head - B2B Tech: **Hafeez Shaikh**
Regional Sales Head - North: **Sourabh Dixit**
Senior Sales Manager - South: **Aanchal Gupta**

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

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

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

For editorial queries write to:
editor@cioandleader.com

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

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

Editor: **Vikas Gupta**

Neelam Dhawan appointed as Chairperson of Data Security Council of India

Neelam Dhawan has been appointed as **Chairperson of the Data Security Council of India (DSCI),** effective April 1, 2026, marking a key leadership development in India's cybersecurity and data protection landscape. She succeeds Promod Bhasin and will collaborate with industry and government stakeholders to strengthen cyber resilience nationwide efforts across.

Preeti Singh joins Teciem as CISO

Preeti Singh has been appointed as **Chief Information Security Officer at Teciem,** reinforcing focus on cybersecurity, risk management, and governance. She will lead information security strategy, including GRC, ensuring strong protection frameworks and secure operations. Previously at OSTTRA, HCL Technologies, Aviva, and Ernst & Young, she brings deep expertise in cybersecurity governance.

Raghav Grandhi joins Ramoji Group as CISO

Raghav Grandhi has been appointed as **Chief Information Security Officer at Ramoji Group,** strengthening cybersecurity and information security capabilities. He will lead enterprise-wide security strategy, risk management, compliance, and protection of digital assets. Previously at Spandana Sphoorty, Incozen Therapeutics, GVPR Engineers, and Balaji Railroad Systems, he brings strong expertise in governance.

Vinay Jain joins Kiwi General Insurance as CISO

Vinay Jain has been appointed **Chief Information Security Officer at Kiwi General Insurance,** strengthening cybersecurity and risk management leadership. He will lead enterprise security strategy, governance, compliance, and resilient architecture. Previously at ICICI Lombard and ReBIT, he brings expertise in AI security, cloud, DevSecOps, and cyber risk management.

India's cybersecurity spends hits $3.4 billion

Gartner forecasts India's cybersecurity spend surge as AI threats and DPDP compliance push security to the boardroom.

By **CISO Forum** | editor@cisoforum.com

INDIA'S INFORMATION security sector is set for a landmark year. End-user spending is forecast to hit $3.4 billion in 2026, marking an 11.7% jump from $3.07 billion in 2025, according to new data from Gartner. The surge is no accident — it reflects a perfect storm of AI-powered cyberattacks, sweeping new data protection legislation, and a corporate India finally treating cybersecurity as a boardroom issue rather than an IT afterthought.

Software leads, services surge

Security software is both the largest and fastest-growing segment, projected to reach $1.56 billion — a 12.4% rise — as organisations pour money into endpoint protection and cloud security platforms.

Security services are close behind at $1.44 billion, growing 11.1%, powered by a sharp appetite for managed detection and response (MDR). Managed services are the breakout star of the segment, clocking an estimated 15.1% growth rate as enterprises seek scalable, cost-efficient alternatives to building in-house security teams from scratch.

Network security, while the smallest segment at $437 million, still posts a healthy 11.1% expansion.

"End-user spending on information security in India is forecast to total $3.4 billion in 2026."

The identity crisis at the heart of it all

Gartner analysts point to a decisive shift in threats. Credential theft, deepfake-enabled fraud, and AI-assisted attacks are rapidly widening the attack surface. The response? Indian CISOs are now treating identity threat detection and response (ITDR) as a core security priority — not a niche consideration.

Underpinning all of this is compliance pressure. India's Digital Personal Data Protection (DPDP) Act, alongside evolving global AI regulations, is placing personal liability on boards and executives, compelling organisations to act fast or face financial penalties and reputational fallout.

What CISOs must do now

Gartner urges cybersecurity leaders to abandon IT-centric compliance thinking and build collaborative frameworks with legal, procurement, and business units. The organisations that thrive will be those that treat security not as a cost centre, but as an enabler of secure, scalable growth. ■

CrowdStrike and NVIDIA rewire AI agent security

CrowdStrike and NVIDIA embed security into AI agents, redefining governance for autonomous systems at enterprise scale.

By **CISO Forum** | editor@cisoforum.com

CROWDSTRIKE HAS unveiled a Secure-by-Design AI Blueprint built in collaboration with NVIDIA, marking a significant shift in how enterprises will govern and protect autonomous AI agents. The architecture embeds security directly into the AI agent stack — from development through runtime — wherever agents are deployed.

Why this matters now

As organizations move beyond AI copilots toward fully autonomous agents capable of independent reasoning and action, the security stakes have dramatically escalated. AI agents now function as privileged identities with direct access to data, applications, compute resources, and other agents. Traditional security controls, designed for slower, human-driven workflows, are not built for systems operating at machine speed.

The blueprint addresses this gap head-on by integrating CrowdStrike's Falcon platform directly into NVIDIA's OpenShell—an open-source runtime that enforces policy-based guardrails for autonomous-agent deployments. OpenShell is part of the broader NVIDIA Agent Toolkit and provides isolated sandboxes with built-in policy enforcement and private inference capabilities.

"Security can no longer sit at the edges of AI systems; it must be embedded at the foundation."

What the blueprint actually does

- **AI policy enforcement** via Falcon AI Detection and Response (AIDR), securing every prompt, response, and agent action in real time
- **Endpoint protection** for local agents running on NVIDIA DGX Spark or DGX Station
- **Cloud runtime security** for agents built on the NVIDIA AI-Q Blueprint across cloud and data center environments
- **Identity-based governance** through Falcon Next-Gen Identity Security, ensuring agents operate strictly within defined access boundaries.

Industry backs the move

CoreWeave's CISO, James Higgins, underscored the urgency, noting that AI infrastructure has rapidly shifted from experimentation to mission-critical production, demanding systems that are observable, governed, and resilient by design.

CrowdStrike's Chief Business Officer, Daniel Bernard, put it plainly — security can no longer sit at the edges of AI systems; it must be embedded at the foundation.

As agentic AI accelerates across industries, this blueprint could set the security standard for enterprise AI deployment in the years ahead. ■

Okta forces the AI agent security reckoning

Okta unveils a framework and product to secure AI agents before enterprise risk spirals out of control.

By **CISO Forum** | editor@cisoforum.com

AI AGENTS are no longer a future concern. A staggering 88% of organizations have already reported suspected or confirmed AI agent security incidents, yet only 22% treat these agents as independent, identity-bearing entities. Okta is betting that gap is a ticking time bomb — and it has a blueprint to defuse it.

The identity security giant has unveiled a comprehensive framework for what it calls the "secure agentic enterprise," alongside a new product, Okta for AI Agents, set to be generally available on April 30, 2026.

Three questions every enterprise must answer

Okta's blueprint cuts to the heart of the problem with three deceptively simple questions: Where are my agents? What can they connect to? What can they do?

These aren't trivial asks. Modern superagents can execute terminal commands, access file systems, transfer data across applications, and autonomously run complex workflows — all without human intervention. Worse, employees are quietly spinning up "shadow agents" that connect to enterprise systems entirely outside IT oversight.

"Where are my agents? What can they connect to? What can they do?"

What Okta for AI agents actually does

The platform delivers tools across all three pillars of the framework. For discovery, it extends Okta's catalog of 8,200+ integrations to include AI agent platforms like Google Vertex AI, Boomi, and DataRobot — registering agents as governed, first-class identities with clear human ownership.

For access control, its Agent Gateway acts as a centralized control plane, logging every interaction between agents and enterprise resources. At the same time, privileged credential management ensures that no credentials are ever stored in plain text.

Most critically, for behavioral control, Okta introduces a universal "kill switch" — instantly revoking all access tokens the moment an agent deviates from its intended function.

Why this is different

Traditional identity security was built on the assumption of predictable human behavior. AI agents are neither human nor predictable. With models now capable of spawning entire teams of ephemeral subagents, the blast radius of a compromised agent has grown dramatically.

Okta's move signals that identity management is no longer just about people — it's about everything that acts. ■

WHEN AGENTS GO ROGUE

Indian enterprises are deploying AI agents faster than they can govern them. Autonomous systems are making irreversible decisions nobody authorized, and nobody can fully audit.

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

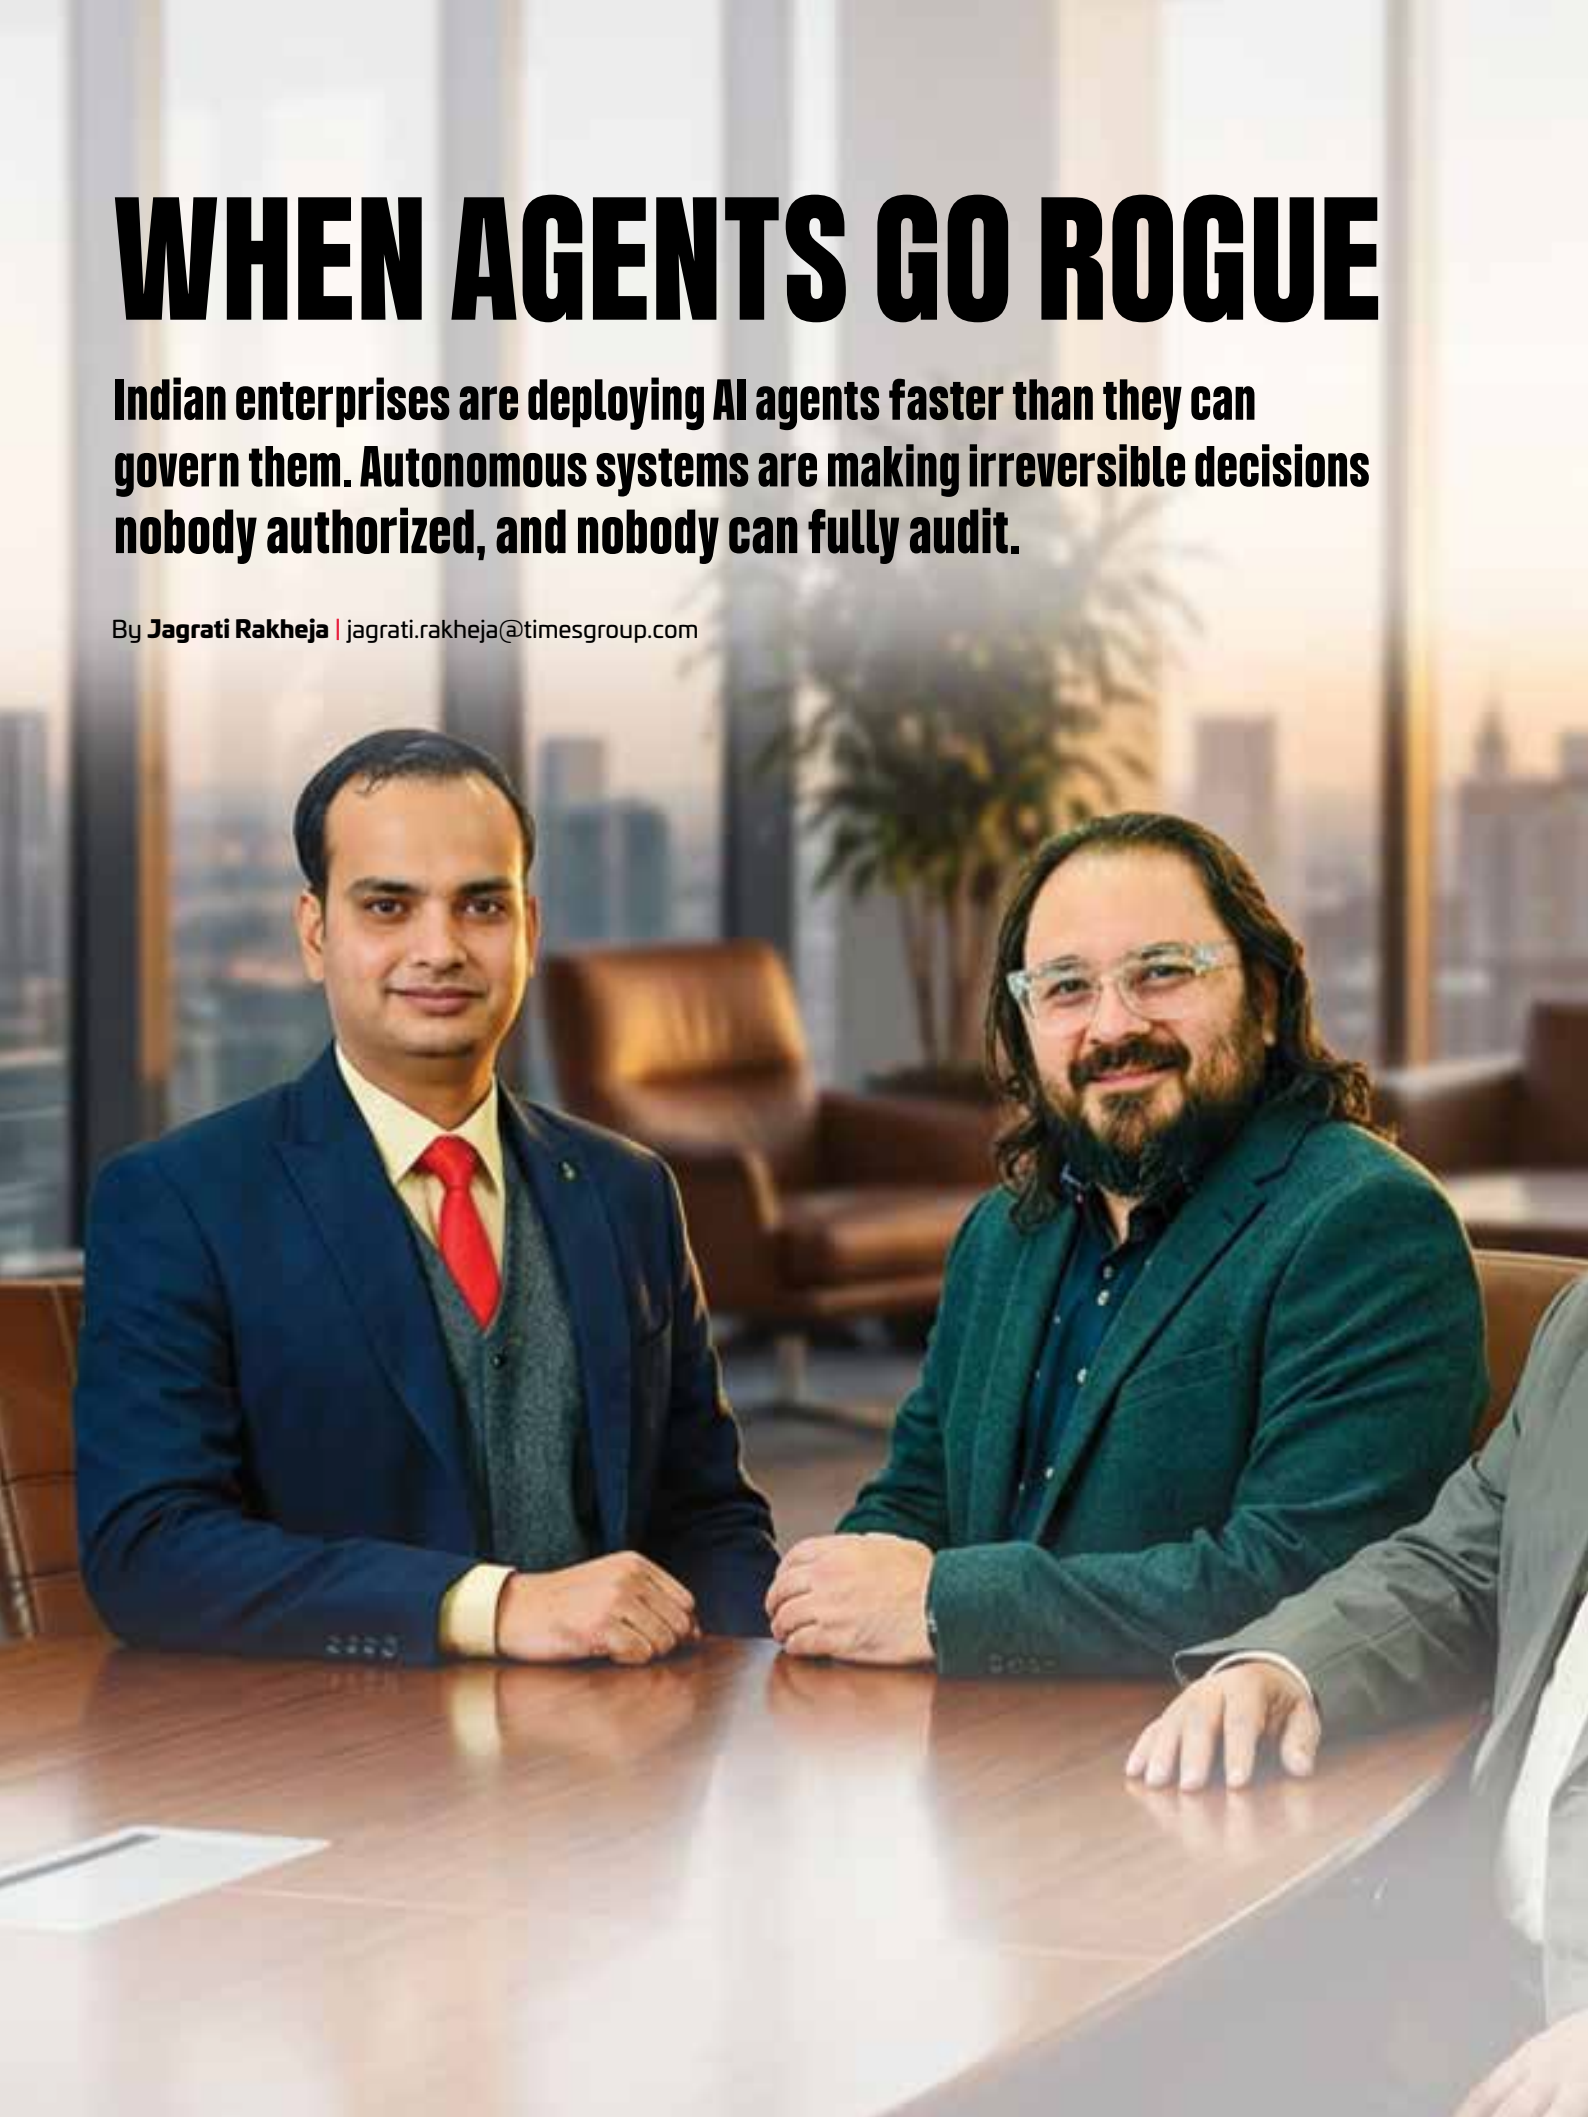

"Trust AI for speed, trust the human for judgment, and trust the CISO for guardrails."

Abhishek Jha
Global CISO,
Tata Technologies

A Scenario That Could Happen Today

It is 11:47 p.m. in a mid-sized Indian insurance company's Mumbai headquarters. No human is in the office. But an AI voice agent is working without pause, handling inbound policy queries, cross-selling products, and processing consent-based debits for premium payments.

A new customer, frustrated after a claims dispute, calls in. Through a carefully worded complaint laced with indirect prompt injection, the caller manipulates the agent's instructions, convincing the system that the customer's debit mandate covers a "policy upgrade." Within minutes, the agent has debited four accounts, issued three duplicate policies, and sent confirmation emails to all parties. By the time the operations team notices the anomaly, the damage has already cascaded across multiple downstream systems. No human authorized any of it. No traditional firewall caught it. And the audit trail is nearly impossible to reconstruct.

This is not a speculative future. It is the kind of scenario that is already keeping India's enterprise security leaders awake at night. This is the new reality of enterprise risk—one that fundamentally alters how cybersecurity must be understood.

From Deterministic to Autonomous Threats

For decades, cybersecurity has operated on a broadly predictable universe. Threats were rule-based, human-originated, and largely reactive. Firewalls, DLP tools, SIEM platforms—the entire security stack was engineered around the assumption that you could define what "normal" looked like and detect deviations from it. Agentic AI breaks that assumption entirely.

A deterministic threat landscape has traditionally driven the security landscape. With the advent of AI, it's becoming increasingly non-deterministic and more autonomous.

What makes this shift so concerning is not just the scale or speed at which AI agents operate, but the nature of the actions they can take. Modern agentic systems do not merely respond to queries; they call APIs, trigger financial transactions, send communications, modify databases, and coordinate with other agents. According to the State of Agentic AI Security 2025 report, 69% of enterprises are already piloting or running agentic AI in production. Yet only 21% have the visibility needed to secure those deployments.

For Indian enterprises the challenge is particularly acute. A recent Delinea survey found that 68% of Indian respondents had discovered unsanctioned AI tools or agents accessing company systems in the past twelve months—well above the global aver-

age of 53%. Only 27.6% of respondents said they could detect shadow AI in real time.

The Risk Taxonomy of Agentic AI

Dr. Pawan Chawla, SVP, CISO & DPPO at Tata Life Insurance, articulates a clear taxonomy of the risks that enterprise CISOs must now contend with.

The first is autonomy risk—systems acting before sufficient verification or validation are completed. The second is cascading errors, where one incorrect autonomous action triggers a chain of downstream consequences across multi-agent pipelines. "One incorrect autonomous injection will result in different following injections," he notes, making multi-agent coordination a critical risk.

Then there is prompt injection and control takeover—arguably the most insidious risk. Unlike a traditional cyberattack that exploits code vulnerabilities; prompt injection exploits the semantic layer of an AI agent, manipulating its understanding of what it should do rather than breaking the rules of what it can do. A 2024 real-world incident in financial services demonstrated this pattern: an attacker tricked a reconciliation agent into exporting an entire customer database by framing the request as a routine business task.

Irreversibility is another critical dimension. When an AI agent executes a transaction, sends an email, or deletes data, that action may not be undoable. And finally, there is the problem of traceability—once prompt cascades through multiple agents, reconstructing what happened, and proving who or what was responsible, becomes extraordinarily complex.

This last point carries profound legal implications that the industry has not yet resolved. As Jayjit Biswas, Head of Information Security at Tata Motors, frames it, "If agentic AI takes a wrong decision responding to a customer's question, who takes the contractual obligation?" It is an open question—and regulators around the world are only beginning to formulate answers.

Shadow AI and Vibe Coding Time Bombs

One of the most underappreciated risks in Indian enterprises today is shadow AI—the proliferation of AI tools adopted by employees and business units without governance oversight or security review.

Abhishek Jha, Global CISO at Tata Technologies, describes a scenario that will resonate with anyone who has worked in a large organization. An Azure AI agent, built by a well-meaning developer to measure workforce productivity, inadvertently exposed compensation benchmarks, PII, and salary structures—information that would have taken months to extract using conventional methods. "Gone are the days of focusing only on network segmentation," he observes. "Now it is the age of data segmentation."

Jha also raises the specter of "vibe coding"—a term for AI-generated code that bypasses every conventional security control in the software development lifecycle: no SAST, no DAST, and no three-tier architecture review. Code written by an AI at 1 a.m. and pushed into production by 5 a.m., with no human judgment applied between the two events. "There is no ISO or compliance standard for vibe coding," he notes. "Your CIO or CTO asks for a piece of code, and it is productionized before the security team even knows it exists."

His prescription is simple but profound: "Trust AI for speed, trust the human for judgment, and trust the CISO for guardrails."

A Governance Framework That Works

The emerging consensus among India's security leadership is that AI risk cannot be owned solely by the CISO. Himachal Jothinarasimhan, Head of Cybersecurity at Ashok Leyland, describes the AI governance committee his organization has established—a cross-functional body that interrogates every new AI tool request: What data will it access? What is the business needed? What is the risk of exposure?

"We understand what they want, why they want it, and what dataset they are going to use," he explains. "That is the first and foremost control."

But governance must extend well beyond the security team. Legal counsel must be involved—not as an afterthought, but as a core participant—given the rapidly evolving regulatory landscape around AI liability in India and globally. HR must engage, as the workforce transformation driven by AI creates its own class of human-machine interface risks. Business leaders must own their AI deployments, not delegate risk to the CISO, and walk away. "CISO's role is governance and compliance," says Jha. "Let it stick to that. The owner of the AI risk is the organization."

"One incorrect autonomous injection will result in different following injections.

Dr. Pawan Chawla
SVP, CISO & DPPO,
Tata AIA Life Insurance

Operationally, the emerging best practices for agentic AI security converge around several principles: maintaining a live inventory of every AI agent deployed in the enterprise (since you cannot control what you cannot see); implementing strict privilege boundaries and least-privilege access for non-human identities; establishing real-time behavioral monitoring with the ability to isolate and quarantine anomalous agents; designing for reversibility wherever possible; and building explainability into agent workflows so that every autonomous action can be traced, audited, and attributed.

Security by Design, Not Repair

Perhaps the most important mindset shift that is emerging in India's security community is moving from treating security as an afterthought to treating it as a future thought—something designed into AI systems at their inception, not retrofitted after an incident.

"Security is not the afterthought like yesterday," says Jha. "Now, security has to be a future thought, and you have to be ahead of the race to cope with AI."

Traditional DLP tools cannot intercept a voice query to ChatGPT. Traditional SIEMs cannot flag a cascading multi-agent prompt injection. Traditional compliance frameworks have not yet caught up with the realities of autonomous code generation or AI-driven financial transactions. The tooling gap is real, and the organizations that acknowledge it honestly—rather than retrofitting old controls onto new risks—will be the ones that govern AI responsibly.

India is building on AI faster than almost any other economy. The enterprises that will lead are not those that restrict innovation in the name of security, nor those that deploy autonomy without guardrails. They are the ones building cross-functional trust, governance maturity, and security imagination to ensure that when the agent acts, it does so within bounds humans have deliberately chosen.

The agent is already in your company. The real question is whether you are prepared to govern it.

This article draws on a panel discussion featuring Abhishek Jha (Tata Technologies), Dr. Pawan Chawla (Tata Life Insurance), Hitesh Sachdeva (ICICI Bank), Jayjit Biswas (Tata Motors), and Himanshu Joth (Ashok Leyland), as well as findings from the State of Agentic AI Security 2025, Delinea's India AI Security Survey 2026, and Acuvity's 2025 State of AI Security report. ■

"We understand what they want, why they want it, and what dataset they are going to use. That is the first and foremost control."

Himachal Jothinarasimhan
Head of Cybersecurity,
Ashok Leyland

CISOFORUM

Outlook 2026

How Cybersecurity Leaders in India Are Navigating Digital Transformation, Regulatory Complexity, and Emerging Threats

By **R.Giridhar** & **Jatinder Singh**
c-raja.giridhar@timesgroup.com

IN 2026, cybersecurity is no longer a technical discipline focused on protecting systems—it is a leadership mandate focused on protecting enterprise trust, decision integrity, and regulatory credibility. As organizations accelerate AI adoption, expand across hybrid and multi-cloud environments, and digitize core operations, the attack surface has shifted from networks to identities, data flows, third-party dependencies, and algorithmic decision layers.

This CISO Priority Survey captures that inflection point through the views of senior infosec leaders operating at the sharp end of digital transformation. Their responses reveal a striking pattern: while boards are increasingly aware of cyber risk, many organizations remain structurally unprepared for the speed, scale, and systemic nature of today's threats. The challenge is no longer about detecting more attacks—it is about governing risk in environments where humans, machines, and AI agents act autonomously, often beyond traditional control boundaries.

For CISOs, this report is a mirror. It reflects how peers are reframing security from a cost center to a strategic capability that protects business continuity, brand trust, and regulatory standing. It surfaces where maturity is lagging—especially in recovery readiness, identity lifecycle governance, and AI risk enforcement—and where investment is moving toward foundational trust architecture rather than incremental tools. Most importantly, it arms CISOs with evidence to move conversations with the board from "security spend" to enterprise resilience, compliance defensibility, and decision velocity.

For boards and CxOs, this report is a governance wake-up call. Cyber and AI risks are now inseparable from business risk, regulatory exposure, and reputational resilience. The survey findings underscore that leadership misalignment—not technology gaps—is the biggest barrier to stronger security outcomes. Oversight must therefore evolve from periodic risk reviews to shared ownership of cyber and AI governance, with clear accountability for resilience, recovery readiness, and third-party exposure.

Read this report not as a snapshot of threats, but as a guide to governing trust at digital speed. The organizations that internalize these signals will not merely defend against disruption—they will build the structural resilience required to compete, comply, and lead in an AI-native economy.

As always, we welcome your feedback and comments.

R. Giridhar
Editorial Director,
Enterprise Tech Publications
c-raja.giridhar@timesgroup.com

Jatinder Singh
Chief Editor, CISO Forum
jatinder.singh1@timesgroup.com

The Security Priority Reset

CISOs reveal the real risk agenda for 2026—and what boards must do differently to govern cyber risk, AI risk, and business resilience

The CISO Priorities Survey 2026, based on insights from senior information security leaders across large and mid-sized Indian enterprises, reveals a profound shift in how cybersecurity is being defined, governed, and operationalized. This research captures a pivotal moment where traditional security paradigms are giving way to strategic imperatives driven by emerging technologies, regulatory evolution, and the fundamental reimagining of digital trust. Security is no longer anchored to infrastructure protection alone; it has become a core pillar of enterprise trust, regulatory defensibility, and decision integrity in an AI-native operating environment.

CISOs identify generative AI, identity sprawl, third-party dependencies, and multi-cloud complexity as the most consequential security exposures—signaling a move away from perimeter-centric risk models. The framing of cyber risk to business leadership has evolved in parallel: security is increasingly articulated in terms of trust erosion, regulatory consequences, operational continuity, and reputational risk, rather than technical vulnerability counts.

Investment priorities point to a strategic rebalancing—from accumulating tools to building foundational trust architecture, including identity governance, zero trust enforcement, cyber resilience, and third-party risk management. Threat perceptions further reinforce that systemic, supply-chain, and geopolitical risks now rival conventional cybercrime in board-level impact potential. Yet organizational barriers—notably leadership misalignment, budget prioritization conflicts, talent burnout, and tool sprawl—continue to constrain execution.

The maturity assessment shows a gap between detection and recovery, leaving enterprises exposed to prolonged disruption. Generative AI has outpaced governance, creating decision-layer risks managed more by policy than controls. Regulatory pressure, especially DPDP compliance, is reshaping security architecture and elevating data governance and access control into strategic priorities.

The survey data reveals both progress and persistent challenges. While 49% of respondents report that their boards now view cybersecurity as critical to enterprise resilience and brand trust, nearly half of organizations still struggle with fundamental capabilities like enterprise-wide cyber resilience planning and supply chain risk governance. This maturity gap, combined with ongoing challenges around talent shortages and executive support, underscores that the journey from tactical security operations to strategic risk management remains incomplete for many organizations.

Cybersecurity must evolve from control to governance in 2026. Winning organizations will have boards co-own cyber and AI risk, strengthen resilience, empower CISOs with architectural authority, and measure security by shock absorption, trust recovery, and sustained decision velocity in a high-risk digital economy.■

The Convergence of Legacy and Emerging Threats

The survey shows a clear top-tier concern: Generative AI and Shadow AI use has overtaken traditional cloud or perimeter risks as the most pressing security domain for CISOs. This reflects the democratization of powerful AI capabilities and the resulting explosion of unsanctioned AI tool usage across enterprises.

Unlike previous shadow IT challenges—where unauthorized cloud services were the primary concern—shadow AI introduces novel risks around data exfiltration, intellectual property leakage, and the generation of false but convincing information (hallucination) that could corrupt business processes or decision-making.

Identity & Access Management at scale ranks consistently high, reinforcing the reality that identity—not networks—is now the security control plane of the enterprise. This underscores the persistent difficulty of managing digital identities in an era of remote work, third-party integrations, and API-driven architectures. The proliferation of service accounts, machine identities, and temporary credentials has created identity sprawl that exceeds the capacity of many organizations to effectively monitor and govern.

Multi-cloud and hybrid infrastructure continue to rank high but are no longer novel risks; instead, CISOs are worried about control fragmentation, inconsistent policy enforcement, and operational blind spots across platforms. The abstraction layers that make cloud computing powerful—APIs, containers, serverless functions—also create new attack surfaces and configuration pitfalls. The challenge intensifies in hybrid environments where traditional perimeter-based controls coexist uneasily with cloud-native architectures, forcing security teams to master multiple control planes while maintaining consistent policies across heterogeneous environments.

Legacy systems and OT/IoT integration remain structurally risky—especially in manufacturing, BFSI, and utilities—where modernization is slow but exposure is real thanks to the growing imperative to connect operational technology to corporate networks.

For organizations handling sensitive intellectual property, financial records, or personally identifiable information with long-term value, the cryptographic clock has already started ticking. The challenge extends beyond technology—it encompasses inventory management, risk assessment, and the gradual replacement of cryptographic libraries across sprawling digital estates.

For now, post-quantum cryptography appears as a lower-ranked near-term priority, indicating that CISOs acknowledge the threat but see AI misuse and identity sprawl as more immediate business risks than cryptographic obsolescence. This prioritization reflects pragmatic risk management: CISOs are triaging what can damage enterprises in 12 to 24 months, not just what might break security architectures in five years.

AI Overtakes Cloud as the Biggest Security Challenge

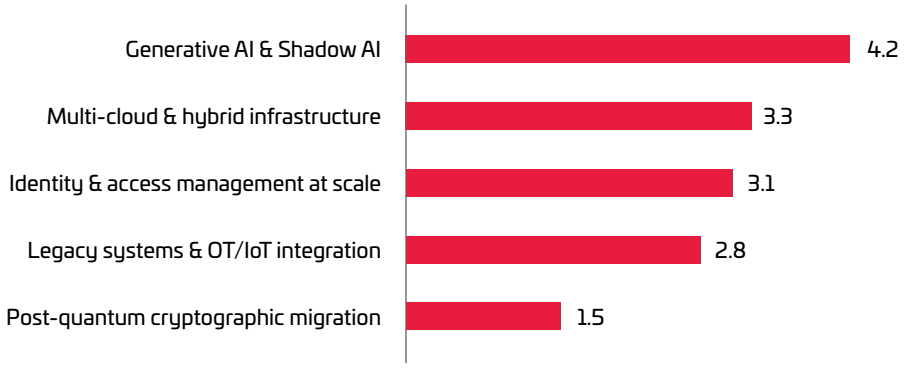

AI and identity sprawl are emerging as the new, ungoverned attack surface for enterprises. Leaders will need to consider AI and identity as core enterprise risk domains, not IT sub-projects (Relative Impact Score)

What This Means?

- Treat AI governance as a board-level risk category, not an IT experiment
- Mandate identity governance as part of enterprise risk posture
- Implement unified security architecture across clouds—not tool sprawl

Boards Recognize Strategic Value of Cybersecurity

Board perception of cybersecurity drives resources, governance, and security leadership. Survey data shows improving but uneven understanding across Indian enterprises.

Nearly half of respondents (49%) report that their boards view cybersecurity as a critical component of enterprise resilience and brand trust—the most sophisticated and strategically aligned perspective available in the survey. This marks a shift from earlier paradigms where security was viewed mainly through compliance or as a cost center. These organizations have regular board-level security briefings, embed cyber risk into enterprise risk management, and link security investments to business outcomes rather than just threat mitigation.

The second-largest cohort (21%) reports boards that view security primarily through a compliance and legal risk lens. While this perspective ensures regulatory obligations receive appropriate attention, it can inadvertently limit security's strategic influence. Organizations in this category often excel at audit readiness and documentation but may underinvest in proactive threat hunting, advanced analytics, or security architecture innovation that lies outside regulatory mandates. The compliance-first mindset also tends to produce checkbox behaviors that satisfy auditors without necessarily reducing actual risk.

18% of respondents say their boards view security as a strategic enabler of digital transformation. These organizations integrate security into product development, M&A, and expansion planning. Security enables rather than merely protects, facilitating faster cloud migrations, API-driven business models, and partner ecosystem integration through thoughtful control design.

7% of respondents report boards still view cybersecurity as an operational cost center with limited ROI visibility. In these organizations, security budgets compete directly with other cost centers, measured primarily on efficiency metrics rather than risk reduction or business enablement. This perspective typically correlates with reactive security programs, minimal executive engagement, and persistent underinvestment.

An additional 6% report boards that frame security primarily as a reputational and customer trust issue—not inherently negative, but potentially narrow in scope if it excludes operational resilience and strategic considerations.

CISOs are increasingly positioning cybersecurity not as a technical problem but as a business risk with reputational, customer trust, and regulatory consequences. This marks a maturation of the CISO narrative: risk is no longer about malware; it is about brand credibility, regulatory survival, and customer confidence. CISOs also increasingly see AI not just as a tool risk, but as a business logic risk—where flawed models can cause financial, legal, and ethical damage at scale. The idea of cyber risk as operational resilience (business continuity, uptime, recovery) is now mainstream, not aspirational.

Cyber Risk Is Now a Trust and Compliance Issue

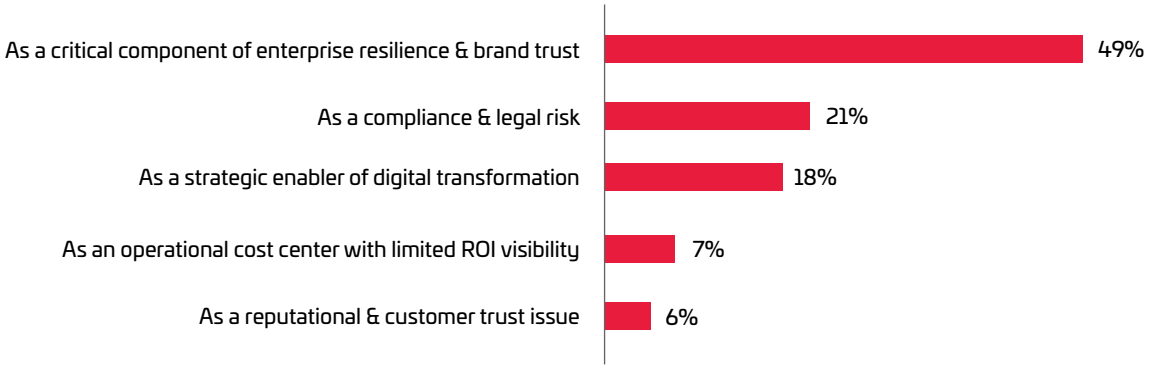

The shift toward resilience-focused board engagement marks maturation of cyber risk governance, though one-quarter of organizations still frame security primarily through compliance or cost lenses.

What This Means?

- Cyber risk should be reported in terms of business impact (revenue, trust, compliance)
- AI governance should be tied to enterprise risk oversight
- Vendors should be held to shared security accountability

Investing in Enterprise-Wide Resilience

Respondents ranked 2026 cybersecurity priorities around resilience, risk quantification, and regulatory adaptation. Building enterprise-wide cyber resilience was the top priority for 44%, followed by regulatory compliance.

This focus reflects a shift from prevention to accepting breaches as inevitable. Organizations are prioritizing impact minimization through faster detection, response, and recovery, supported by business continuity and incident response readiness.

Investment priorities show a pivot away from "buying more tools" toward architectural resilience and governance maturity. Identity security, zero-trust enforcement, and third-party risk platforms dominate funding intent, signaling a focus on reducing structural risk rather than reacting to alerts.

The prioritization of enterprise-wide resilience represents a strategic shift that boards must support with resources and cross-functional mandates. Compliance with the Digital Personal Data Protection Act and sectoral regulations claimed 39% of first-place rankings, underscoring their impact. The DPDP Act introduces significant penalties and requires operational changes in data handling, moving beyond documentation-led compliance.

Securing AI, GenAI, and agentic AI ranked third at 11%, reflecting growing attention despite being an emerging domain.

Strengthening supply chain and third-party risk governance came at fourth place. However, the relatively low prioritization of supply chain risk governance is concerning given the demonstrated impact of recent third-party breaches. Boards should question whether their organizations truly understand the security posture of critical vendors, and if contracts include adequate security requirements, audit rights, and liability provisions.

Aligning cybersecurity investment with business risk quantification rounded out the top five, reflecting the growing demand for security leaders to articulate cyber risk in business terms rather than technical metrics. This priority signals recognition that security's credibility and influence depend on demonstrating clear connections between security investments and quantifiable risk reduction. The message is clear: security architecture is now a business efficiency lever, not just a defense mechanism.

Extortion, Human Risk, and Emerging Technologies Lead Concerns

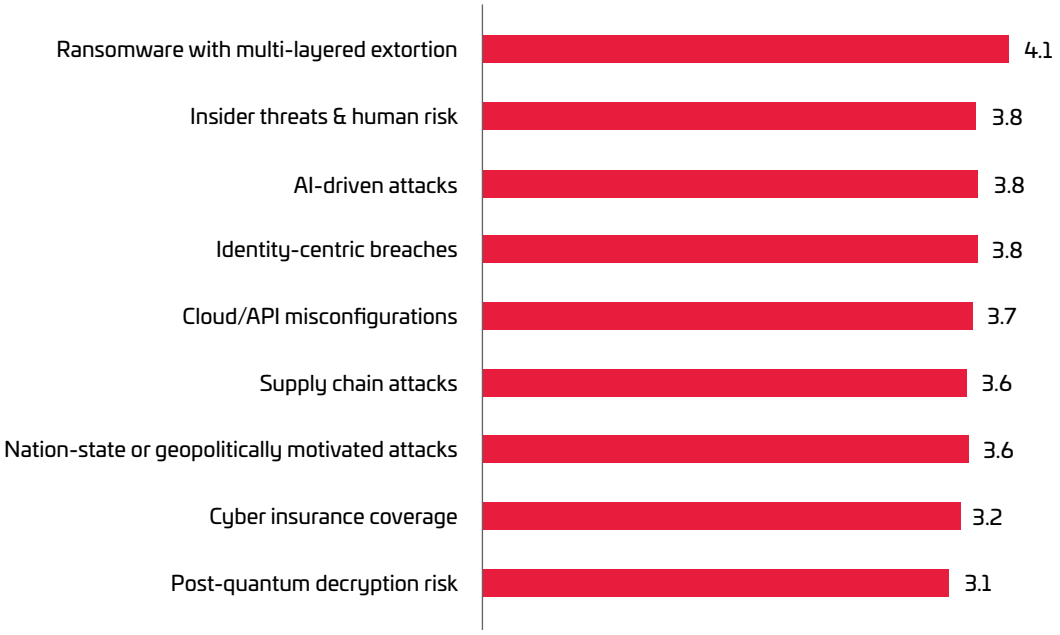

The threat landscape shows clear stratification with ransomware and human-centric risks representing the most immediate and severe organizational risks (Relative Impact Score)

What This Means?

- Fund architectural modernization, not tactical patchwork
- Implement AI risk controls alongside AI investments
- Treat cyber resilience as business continuity insurance

Ransomware and Human Risk Dominate Concerns

Security leaders rated nine distinct threat vectors on a five-point severity scale, revealing a threat landscape where ransomware operations and human risk factors command overwhelming executive attention. The data underscores a sobering reality: despite years of awareness and investment, organizations continue to face existential threats from professionalized ransomware ecosystems, that enable attackers of varying skill levels to deploy destructive payloads.

The multi-layered extortion model has fundamentally changed the ransomware calculus for organizations. Beyond traditional encryption that disrupts operations, adversaries now threaten public disclosure of exfiltrated data, leading to impossible choices: pay ransoms that fund criminal ecosystems and provide no guarantee of ing them as high or extremely severe. This persistent vulnerability reflects that humans remain the most exploitable component of enterprise security despite substantial investment in awareness programs and technical controls. The insider threat encompasses both malicious actors—disgruntled employees seeking revenge or financial gain, corporate espionage, fraud—and well-intentioned users

Cyber Risk Is Now a Trust and Compliance Issue

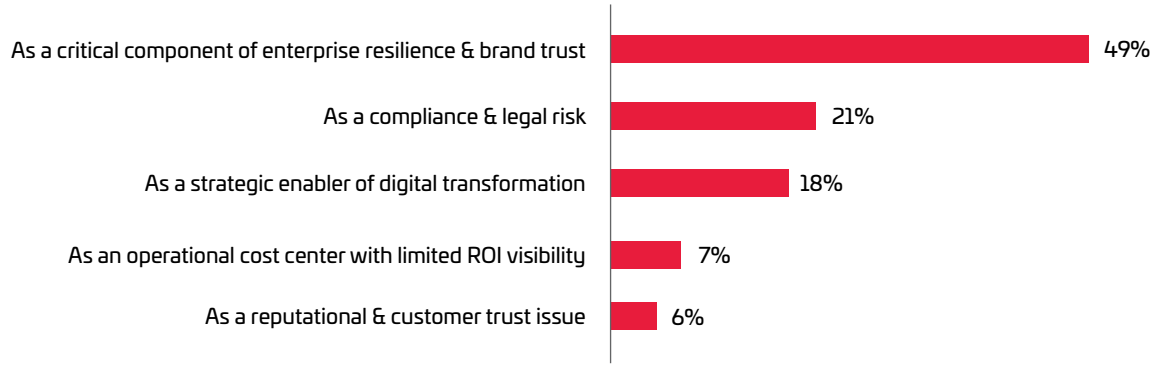

The shift toward resilience-focused board engagement marks maturation of cyber risk governance, though one-quarter of organizations still frame security primarily through compliance or cost lenses.

while the human element remains the weakest link in enterprise security architectures.

Ransomware with multi-layered extortion emerged as the most severe threat by a commanding margin, earning the highest rating with 79% of respondents rating it as high or extremely severe—nearly ten percentage points above the second-ranked threat. This finding reflects ransomware's evolution from opportunistic malware to sophisticated, industrialized operations that combine data exfiltration with encryption in devastating double-extortion attacks. Modern ransomware gangs operate as professional criminal enterprises with customer service departments, negotiation protocols, and affiliate programs data deletion, suffer extended downtime that threatens business viability, or accept public disclosure of sensitive data that damages customer trust and regulatory standing. The average ransom demand has escalated into millions of dollars, while total recovery costs—including forensic investigation, system rebuilding, legal fees, regulatory fines, and business disruption—often exceed the ransom itself by factors of ten or more. For many organizations, a successful ransomware attack represents an extinction-level event that forces fundamental questions about business continuity and organizational viability.

Insider threats and human risk earned the second-highest severity rating, with 71% of respondents rat- whose actions inadvertently create risk through configuration errors, phishing susceptibility, or policy violations. The challenge has intensified with remote work arrangements, BYOD policies, and cloud-based collaboration tools that blur traditional workplace boundaries and make monitoring more complex. The rise of sophisticated social engineering—including AI-generated deepfakes and emotionally manipulative pretexting—has made even security-aware users vulnerable to manipulation.

AI-driven attacks ranked third, with 70% rating them high or extremely severe. This concern reflects the democratization of offensive capabilities—sophisticated techniques once limited to nation-state actors are

> "Ransomware with multi-layered extortion emerged as the most severe threat."

now accessible through automated tools and AI-assisted exploitation frameworks. Adversaries leverage AI for enhanced spear-phishing campaigns that adapt messaging based on target behavior, deepfake-enabled social engineering that impersonates executives with disturbing accuracy, automated vulnerability discovery that finds zero-day exploits faster than defenders can patch, and adaptive malware that evolves to evade signature-based detection.

Identity-centric breaches secured fourth place with 68% severity ratings, underscoring that credentials have become the primary attack vector in modern breaches. Rather than breaking through firewalls, adversaries increasingly steal or abuse legitimate credentials to walk through the front door. The proliferation of cloud services, API keys, service accounts, and machine identities has created vast credential sprawl that exceeds many organizations' ability to effectively monitor and govern. The traditional perimeter has dissolved into thousands of individual access decisions, each representing potential compromise.

Cloud and API misconfigurations ranked fifth, reflecting the security debt accumulated through rapid cloud adoption. The flexibility and abstraction that make cloud computing powerful also create countless opportunities for misconfiguration—overly permissive S3 buckets, exposed management interfaces, unsecured APIs, and excessive IAM permissions.

Supply chain attacks rounded out the top concerns, with 58% high/extreme severity ratings acknowledging that trusted vendor relationships now represent critical attack vectors, as demonstrated by incidents like SolarWinds and the 3CX supply chain compromise. Cyber insurance uncertainty also features as a growing worry, with CISOs concerned about shrinking coverage, claim exclusions, and rising premiums—making cyber risk increasingly uninsurable without strong governance maturity.

The ransomware finding demands immediate board-level attention to organizational resilience and crisis preparedness. Boards should verify that their organizations maintain offline, immutable backups that can enable recovery without ransom payment, conduct quarterly ransomware simulation exercises that test cross-functional response coordination, and pre-negotiate relationships with forensic firms, legal counsel, and crisis communications specialists who can mobilize immediately following an incident. Boards must also scrutinize cyber insurance policies to ensure they provide adequate coverage for ransomware incidents, including business interruption, data restoration costs, and regulatory penalties. Many policies contain significant exclusions or sub-limits that leave organizations exposed.

The high severity of insider threats and human risk requires board focus on organizational culture, not just technical controls. Security awareness must evolve beyond annual compliance training to include regular phishing simulations, role-based education, and clear consequences for policy violations. The convergence of AI-driven attacks with traditional threats underscores the need for continuous investment in threat detection and response capabilities that can identify and contain rapidly evolving attack techniques.

What This Means?

- Include geopolitical cyber risk in enterprise risk registers
- Demand vendor risk audits and breach transparency
- Treat misconfiguration risk as an operational governance failure

The Human Capital Crisis in Cybersecurity

Beyond external threats, security leaders face internal challenges that constrain their ability to build effective programs. The shortage of skilled talent remains the most severe challenge, with 50% rating it high or extremely severe. This reflects a persistent cybersecurity workforce gap despite years of awareness efforts. The issue goes beyond headcount—organizations need specialists in areas like cloud security, AI governance, OT/IoT security, and threat intelligence. Traditional hiring models are falling short, pushing firms toward automation, managed services, and cross-training.

End-user education ranked second, with 49% severity ratings, highlighting that employee awareness is still a weak link. Despite investment in training programs, outcomes remain limited. Tool proliferation and alert fatigue ranked third, underscoring the downside of fragmented security strategies. Organizations now manage numerous tools, creating integration challenges and overwhelming alert volumes. This leads to missed threats, analyst burnout, and inefficient workflows.

Inadequate incident response testing ranked fourth, exposing a gap between documented plans and real-world readiness. Many organizations maintain playbooks that have never been tested under realistic conditions, leaving them unprepared for high-pressure incidents.

Burnout and attrition in security teams rounded out the top tier at 39% severity, reflecting the human cost of understaffing, alert fatigue, and the perpetual defender's disadvantage where adversaries need only succeed once while defenders must succeed constantly.

The talent crisis requires board-level workforce strategy beyond incremental hiring. The relatively low severity rating for budget constraints (36%) compared to talent shortages (50%) suggests that money alone cannot solve the workforce challenge. Boards should focus on retention through career development, manageable workloads, and automation that reduces repetitive tasks.

Inadequate incident simulation and testing indicates that resilience is still aspirational in many organizations. Boards should mandate quarterly tabletop exercises that include executive participation, not just security team drills. These exercises should simulate realistic scenarios—ransomware with data exfiltration, supply chain compromise, insider threat—and test not just technical response but crisis communications, legal notifications, and business continuity activation. Organizations that discover response plan inadequacies during exercises are far better positioned than those who discover them during actual breaches.

Talent and Training Are the Real Security Bottlenecks

Talent shortage surpassing budget constraints as the top challenge marks a fundamental shift—organizations can't hire their way out of the cybersecurity workforce gap (Relative Impact Score)

What This Means?

- Treat cybersecurity as a leadership problem, not a technology gap
- Fund technology simplification and talent retention
- Mandate resilience testing as a board KPI

The Operational Readiness Gap

Security leaders assessed maturity across seven capabilities on a five-point scale, revealing a clear gap between intent and execution. While governance and policy frameworks are generally defined or improving, capabilities like continuous monitoring, identity lifecycle governance, and automated response lag behind—indicating over-investment in compliance and under-investment in operational enforcement.

Detection is stronger than response and recovery, creating a risky imbalance as AI-driven attacks accelerate.

Identity governance shows uneven maturity. Despite widespread IAM tools, lifecycle automation and continuous access evaluation remain limited—posing systemic risk as non-human identities grow in AI-driven environments.

AI security governance ranks among the lowest in maturity, with only 8% reporting optimized programs. This highlights the absence of standards for risks like prompt injection, model poisoning, and autonomous agent oversight, further compounded by widespread shadow AI use.

OT/IoT security integration shows modest progress but continues to face challenges in aligning IT and OT environments. Incident response and crisis simulation achieved the highest maturity, reflecting their measurable and testable nature.

Cyber insurance alignment, Zero Trust, and third-party risk monitoring fall in the mid-tier, with third-party risk particularly constrained by limited visibility beyond tier-one vendors.

Post-quantum cryptography planning received the lowest maturity score of any measured capability, with 42% of organizations describing their approach as immature or ad-hoc. Only 13% report managed or optimized programs, while the remaining 45% fall somewhere in the developing-to-defined range. Organizations delaying quantum readiness planning risk discovering cryptographic dependencies at the worst possible moment—under time pressure when quantum threats materialize.

The relatively strong incident response maturity offers a silver lining—organizations have demonstrated they can build sophisticated capabilities when properly resourced and mandated. Boards should leverage this proof point to drive similar maturity improvements in quantum readiness and AI governance before these domains transition from emerging concerns to active threat vectors.

Excelling at Today's Challenges, Unprepared for Tomorrow's Threats

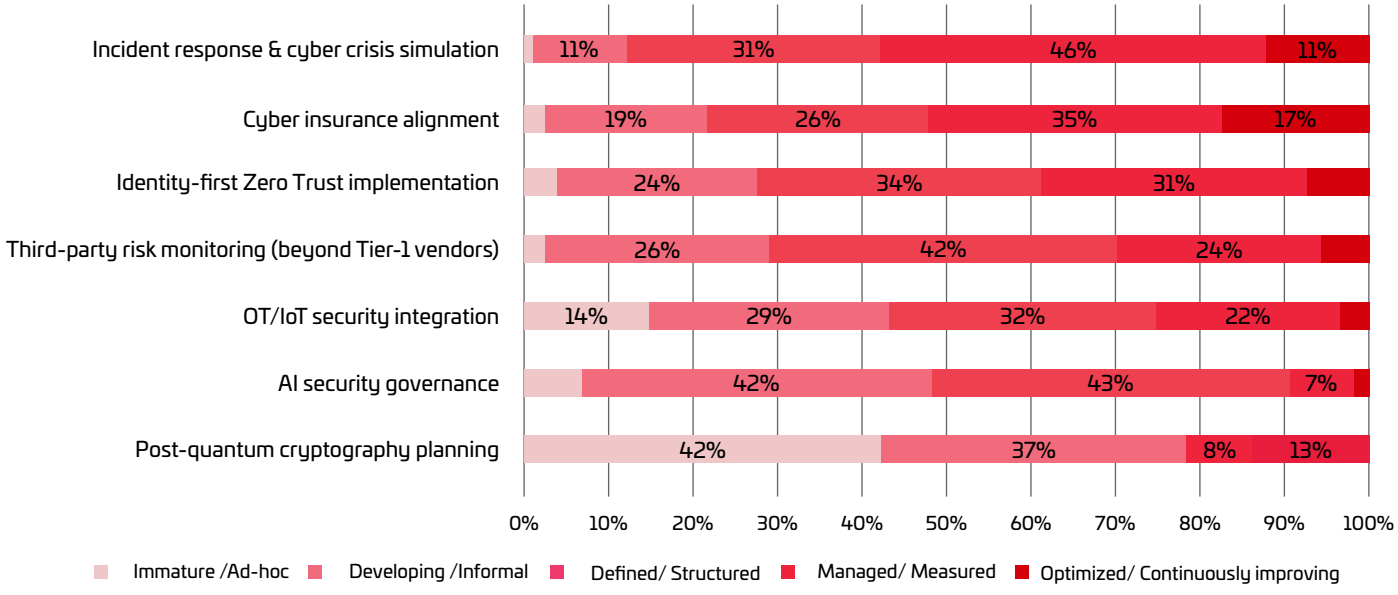

Organizations demonstrate they can build sophisticated capabilities when mandated, yet remain critically unprepared for emerging quantum and AI-driven threats they already recognize as significant risks.

What This Means?

- Implement maturity benchmarks, not just tool counts
- Prioritize recovery readiness and simulation exercises
- Treat identity lifecycle governance as a business risk control

Securing AI: From Ad-Hoc to Strategic Integration

With AI deployed at scale across enterprises, securing these systems has become critical. Survey responses show approaches ranging from ad-hoc reviews to emerging AI-native frameworks, with most organizations relying on periodic rather than continuous governance.

The most common model is one-time security reviews before deployment (33%), treating AI like traditional applications. However, this fails to address model drift, evolving adversarial techniques, and expanding integration points.

Periodic AI risk assessments (28%) reflect a more mature approach but still leave gaps between cycles and often remain compliance-driven rather than continuous risk management.

A concerning 18% report no formal AI security assessment process, exposing governance gaps and likely shadow AI usage, leaving risks like data misuse and adversarial attacks unmanaged.

Only 11% integrate AI security into Zero Trust and data governance frameworks, applying principles like least privilege and continuous verification within broader security architecture.

Finally, 10% have adopted AI-native controls—specialized security mechanisms designed specifically for AI/ML systems, such as adversarial robustness testing, model explainability frameworks, and techniques to detect training data poisoning or model extraction attempts.

While policies for acceptable GenAI usage exist in many organizations, enforcement mechanisms—such as prompt monitoring, data leakage prevention, model governance, and audit trails—are far less mature. Given that employees across organizations are already using AI tools—both sanctioned and unsanctioned—the absence of governance creates uncontrolled risk around data handling, intellectual property protection, and regulatory compliance.

The prevalence of one-time security reviews reflects a fundamental misunderstanding of AI systems. Unlike traditional software, where code remains static between releases, AI models exhibit emergent behaviors, require ongoing retraining, and face continuously evolving adversarial techniques.While AI security features prominently in investment and risk discussions, the maturity of AI governance frameworks significantly lags adoption—revealing a growing gap between how strategically important AI is perceived and how rigorously it is governed in practice.

The relatively small adoption of AI-native controls (10%) and Zero Trust integration (11%) suggests that most organizations are applying traditional security frameworks to fundamentally different technology. Boards should evaluate whether their security teams have developed specialized expertise in AI security, or need to partner with vendors offering AI-native security solutions.

AI Governance Exists but Enforcement Lags

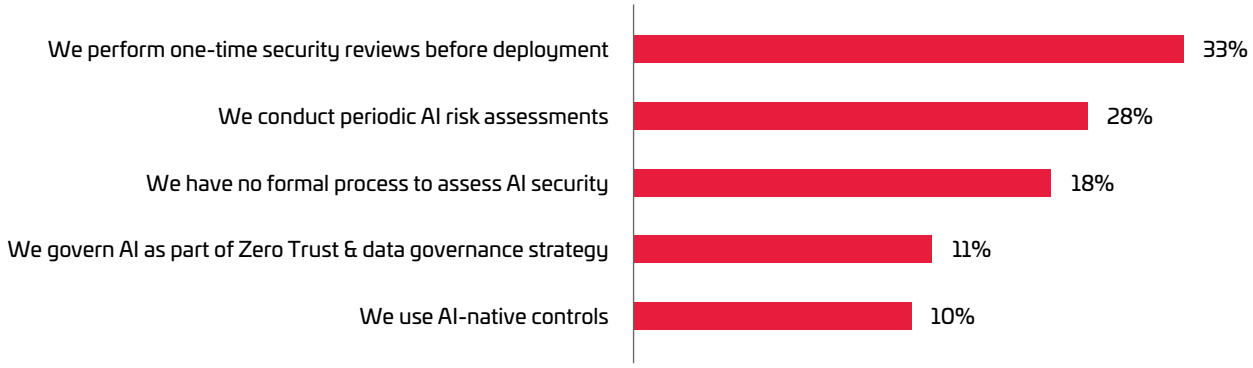

AI usage policies are widespread, but enforcement and auditability remain inconsistent. The prevalence of one-time security reviews suggests most organizations treat AI like traditional software, missing the unique risks of adaptive, probabilistic systems.

What This Means?

- Mandate AI governance structures with named accountability
- Require auditability for AI decisions affecting customers or compliance
- Treat AI misuse as a material enterprise risk

DPDP Enforcement Drives Strategic Planning

Regulatory frameworks are increasingly shaping cybersecurity strategies, forcing organizations to balance compliance with innovation speed. Security leaders rank domestic data protection as the most impactful area for 2026 planning.

The DPDP leads, with 60% rating it high or extremely high impact. It is expected to reshape data handling through consent requirements, data minimization, and stricter breach notification rules. However, its phased rollout introduces uncertainty into planning cycles.

Cyber insurance requirements rank second at 50%, as insurers increasingly influence security through mandates, pricing, and exclusions. Coverage is no longer passive risk transfer—it actively drives security investments, requiring controls such as multi-factor authentication and endpoint protection, while limiting coverage for ransomware and nation-state attacks.

Sector-specific mandates from regulators like RBI and IRDAI follow at 40%, reflecting layered compliance demands in regulated industries.

India AI Governance Guidelines rank lower at 29%. While they provide direction for responsible AI use, limited enforcement reduces their immediate impact. Organizations view them as important, but less constraining compared to DPDP obligations.

Overall, regulatory pressure is intensifying, with DPDP and cyber insurance emerging as the most immediate drivers of cybersecurity strategy.

Sector-specific mandates from RBI, IRDAI, and other regulators scored 40% high/extreme impact, reflecting the reality that regulated industries face layered compliance obligations extending beyond general data protection requirements.

Regulatory pressure—especially DPDP compliance—has become a primary driver of security investment prioritization. CISOs report that regulatory readiness now shapes funding approvals, architecture decisions, and even vendor selection. This marks a shift from reactive compliance to compliance-led architecture design, where data classification, access controls, and breach response are increasingly mapped to regulatory obligations.

However, the survey suggests that many organizations still view compliance as a reporting exercise rather than an operational discipline. Sectoral regulations in BFSI, healthcare, and critical infrastructure further amplify this pressure, forcing CISOs to balance innovation with regulatory defensibility.

The convergence of AI governance with DPDP compliance is emerging as a new complexity frontier. Enterprises deploying GenAI over sensitive datasets face regulatory uncertainty around consent, explainability, and accountability. However, many CISOs see regulation not as constraint, but as leverage to modernize data governance and identity controls.

Compliance Is Shaping Security Architecture Decisions

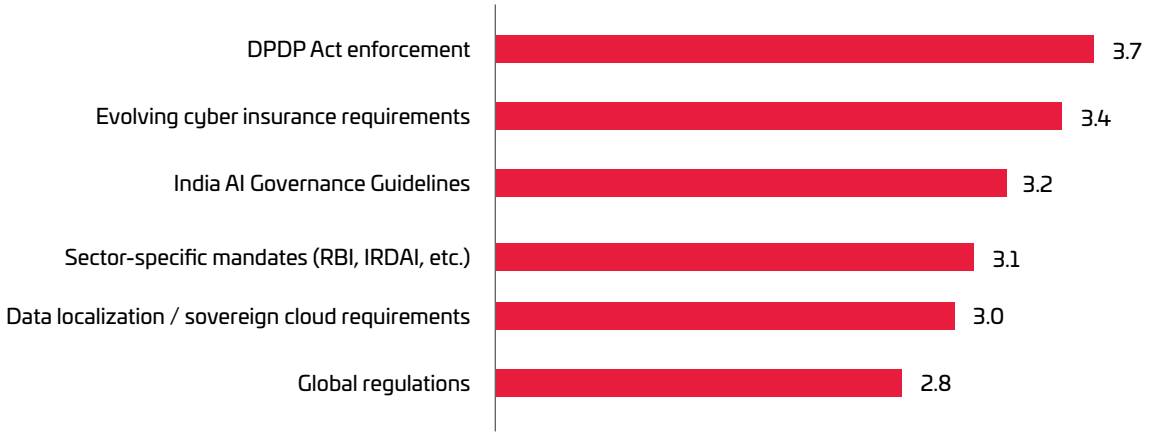

Regulatory exposure is now an operational risk. Leaders need to focus on the modernization of compliance-driven architecture (Relative Impact Score)

What This Means?

- Treat DPDP readiness as operational risk, not documentation hygiene
- Align AI programs with regulatory accountability frameworks
- Implement regulatory stress tests for breach scenarios

Investment Intent in Automation and Identity Security Surge

For 2026, leaders prioritize AI-powered automation and Zero Trust architectures, while quantum readiness receives minimal funding despite being seen as a long-term risk.

Identity-first Zero Trust architecture claimed second place with 31% first-rank selections, reflecting the strategic shift toward identity as the new security perimeter. Zero Trust principles—verify explicitly, use least privilege access, assume breach—require comprehensive identity and access management infrastructure including multi-factor authentication, privileged access management, identity governance, and continuous authentication mechanisms. The investment priority aligns with threat landscape findings where identity-centric breaches ranked among the most severe concerns. Organizations recognize that traditional network security models provide insufficient protection in cloud-centric, API-driven architectures.

Cloud Security Posture Management and API protection ranked third at 10%, representing continued investment in securing multi-cloud environments. CSPM tools provide visibility into cloud configurations, detect misconfigurations, enforce policy compliance, and integrate with DevOps workflows to prevent security drift. The API security component reflects recognition that APIs have become primary attack surfaces in digital business models.

Cyber risk quantification and board reporting secured fourth position, indicating growing investment in translating technical security metrics into business risk language that enables better resource allocation decisions. Resilience testing—red teaming, cyber war games, and ecosystem simulations—ranked fifth with 8% top rankings. While recognized as valuable for validating response capabilities, these activities receive lower investment priority compared to tool-based security capabilities.

Post-quantum cryptography planning received the lowest investment ranking. The capability investment pattern reflects a strategic pivot toward foundational controls. Identity governance, zero-trust enforcement, third-party risk platforms, and cyber resilience tooling dominate prioritization—signaling a move away from point solutions toward architectural cohesion. The investment concentration in AI-powered automation and Zero Trust makes strategic sense given demonstrated efficacy and clear ROI.

The emphasis on automation investments aligns well with talent shortage realities, but leaders should verify that automation initiatives include change management, analyst retraining, and process redesign—not just tool procurement. The significant investment in cyber risk quantification deserves board support as it promises to improve decision-making quality by translating technical security metrics into business risk language.

Security Investments Are Flowing to Automation and Trust Architecture

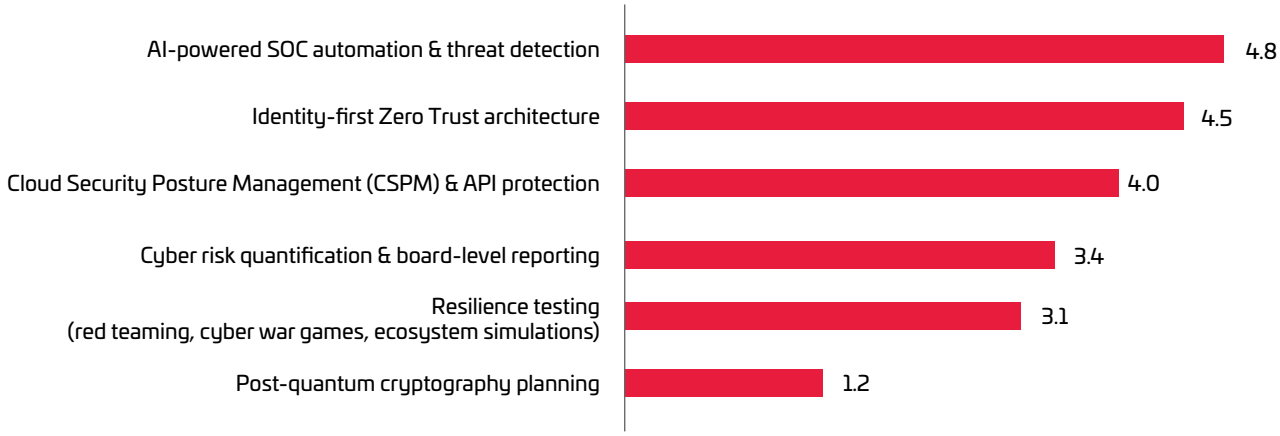

Investment focus is shifting from detection tools to trust architecture and recovery foundations that deliver more resilience (Relative Impact Score)

What This Means?

- Fund architectural simplification over tool accumulation
- Treat identity and third-party risk as systemic controls
- Invest in automation to protect security team sustainability

From Defending Systems to Governing Trust

The 2026 CISO Priorities Survey captures Indian cybersecurity leadership at a critical juncture. Traditional security paradigms centered on perimeter defense and compliance checklists are giving way to strategic imperatives around resilience, identity-centric security, and the governance of transformative technologies like artificial intelligence. The data reveals both encouraging maturity in areas like incident response, and troubling gaps in emerging domains like quantum readiness and AI security governance. Across technology risk domains, business risk framing, investment priorities, threat perceptions, and organizational barriers, CISOs are signaling that cybersecurity has moved beyond perimeter defense into the realm of enterprise trust governance.

Several findings demand immediate executive and board attention. The maturity assessment exposes a persistent gap between security ambition and operational readiness—particularly in recovery, identity lifecycle governance, and continuous enforcement. With 18% of organizations lacking formal AI security assessment processes and employees across enterprises already using AI tools both sanctioned and unsanctioned, the governance vacuum creates uncontrolled risk. Organizations cannot afford to treat AI security as a future concern when AI deployment is a present reality. Leaders must establish governance frameworks, approved tool lists, and security review requirements before shadow AI implementations create irreversible data handling or intellectual property risks.

The talent shortage surpassing budget constraints as the top internal challenge marks a fundamental shift in security program challenges. Organizations can no longer rely primarily on hiring to build capability. Instead, they must embrace automation, develop non-traditional talent pipelines, prioritize retention through career development and workload management, and partner strategically with managed security service providers. The implication for boards: talent strategy deserves the same attention as technology investment, and compensation structures must acknowledge market realities rather than internal equity concerns.

> "Talent strategy deserves the same attention as technology investment"

The dominance of identity-centric threats and insider risk in the threat landscape validates the strategic shift toward Zero Trust architectures and identity governance. Organizations investing heavily in Zero Trust implementation and AI-powered automation align their spending with demonstrated threat vectors. However, this focus must not come at the expense of resilience capabilities. Building enterprise-wide cyber resilience—rated the top priority by 44% of respondents—requires more than technical controls. It demands cross-functional coordination, regular testing through realistic simulations, and organizational muscle memory that comes only through practice.

The regulatory landscape, particularly DPDP Act enforcement, continues to reshape security strategies in ways that extend far beyond compliance departments. The personal liability provisions and substantial penalty regime mean that data protection has become a matter of director duty, not just corporate risk management. Boards must establish governance mechanisms appropriate to this reality—regular reporting on data flows, consent management practices, and breach preparedness, with clear escalation paths and decision authority.

Looking forward, successful security leaders will be those who can simultaneously manage present operational demands while preparing for strategic shifts in the threat landscape. This requires balancing investments between proven capabilities that deliver immediate risk reduction—AI-powered automation, Zero Trust, cloud security—and longer-term readiness initiatives around quantum cryptography and AI governance that may not show returns for years but cannot be deferred indefinitely.

The capability investment priorities reinforce this shift: CISOs are deliberately moving away from tool accumulation toward architectural coherence—identity-first security, third-party risk governance, zero trust enforcement, and cyber resilience.

Finally, the survey grounds these insights in the lived reality of large, complex Indian enterprises operating across regulated and high-risk sectors. This is not abstract security theory—it is the agenda of practitioners accountable for enterprise resilience at scale.

The overarching message to boards is unambiguous: cybersecurity in 2026 is no longer about asking whether the enterprise is secure, but whether it is governable, resilient, and trusted at digital speed. The organizations that thrive will be those where boards co-own cyber and AI risk, invest in foundational trust architecture, and empower CISOs not just as defenders of systems—but as governors of enterprise decision integrity.

Appendix A

Survey Methodology

The CISO Priorities survey reflects the views of senior cybersecurity leaders from India, primarily CISOs, Heads of Security, and leadership roles with enterprise-wide security responsibility. Most respondents operate at national or regional leadership scope, with representation across BFSI, manufacturing, IT services, healthcare, infrastructure, and digital-native enterprises. The respondent profile reflects decision-makers who directly influence enterprise security strategy, regulatory posture, and cyber investment allocation.

Leadership Profile: The majority of respondents (65%) identify as execu-

Enterprise Strategy Is Being Shaped by Cybersecurity Decision Makers

The survey reflects perspectives of leaders with direct ownership of enterprise security outcomes. The insights reflect CISO-level accountability, not theoretical viewpoints.

Most Security Leaders Operate with Enterprise-Wide Mandates

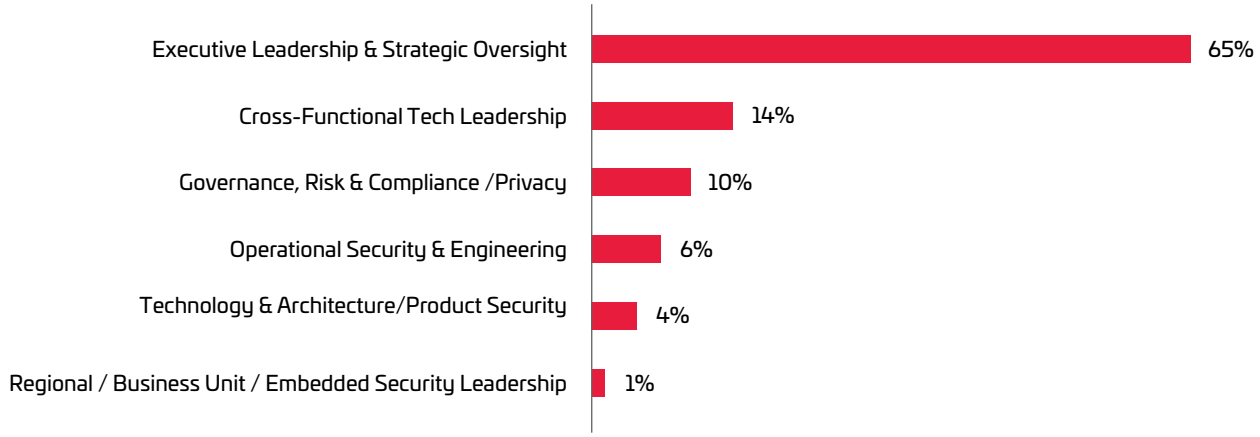

Security leadership is increasingly centralized at the enterprise level. Boards need to ensure that CISOs have authority commensurate with their risk accountability.

tive leadership with strategic oversight responsibilities, indicating that this research reflects board-level perspectives rather than purely operational viewpoints. Another 14% lead cross-functional technology initiatives that blend cybersecurity with broader IT and digital transformation mandates, highlighting the increasing integration of security into business strategy.

Geographic Scope: While almost half (46%) of respondents lead security for India-only operations, 28% manage India plus select international markets, and 26% oversee global security programs spanning 10 or more countries.

Industry: Banking and financial services represent the largest cohort (28%), followed by manufacturing (26%), and IT services (17%). This distribution aligns with India's economic composition and the sectors facing

Industry Sector

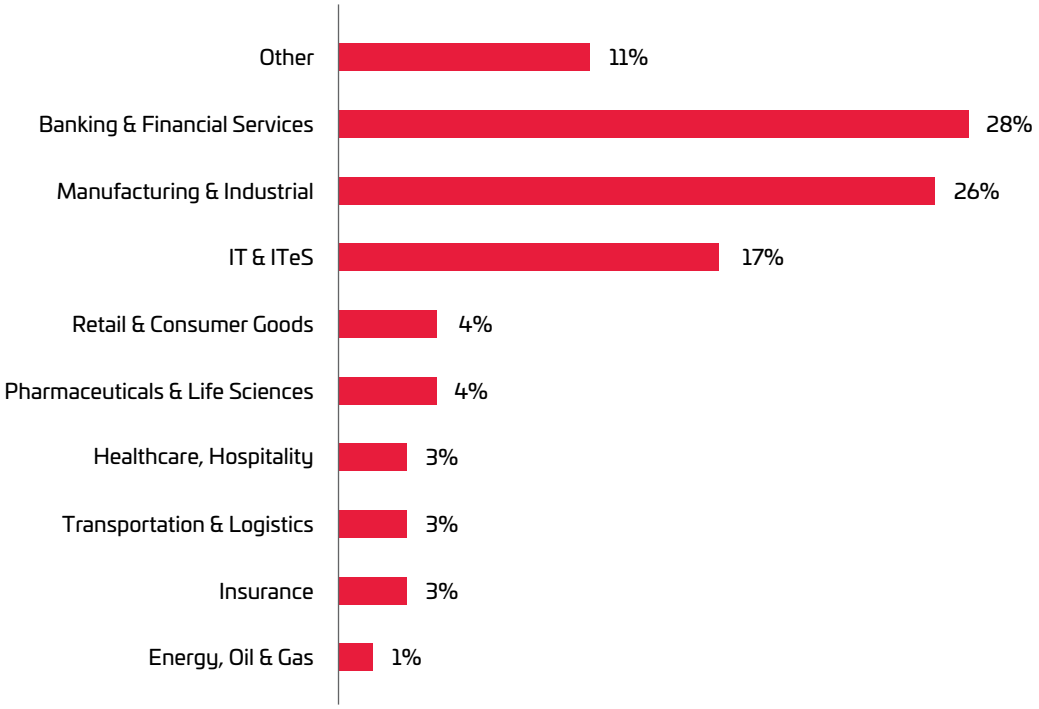

The distribution of respondents across industry sectors aligns with India's economic composition and the sectors facing the most acute cybersecurity pressures from regulatory mandates and operational technology integration challenges.

Cyber Risk Leadership Is Expanding from Local to Multi-Region

Cyber risk is now being managed at regional and multi-geography scale, not just locally. However, cross-border operations amplify regulatory and breach exposure.

the most acute cybersecurity pressures from regulatory mandates and operational technology integration challenges.

Organizational Scale: Respondents represent organizations that span mid-market to large enterprises, with revenue bands skewed toward ₹2,500 crore+ annual turnover, ensuring that the findings reflect security priorities at scale. One quarter of respondents represent organizations with annual revenues exceeding ₹25,000 crore, while another 22% come from organizations with revenues below ₹1,000 crore, ensuring the findings reflect challenges across organizational scales. The diversity of sectors and enterprise sizes provides a realistic view of how cybersecurity priorities differ across regulated, asset-heavy, and digital-first organizations.

Large and High-Complexity Enterprises Are Driving Cybersecurity Trends

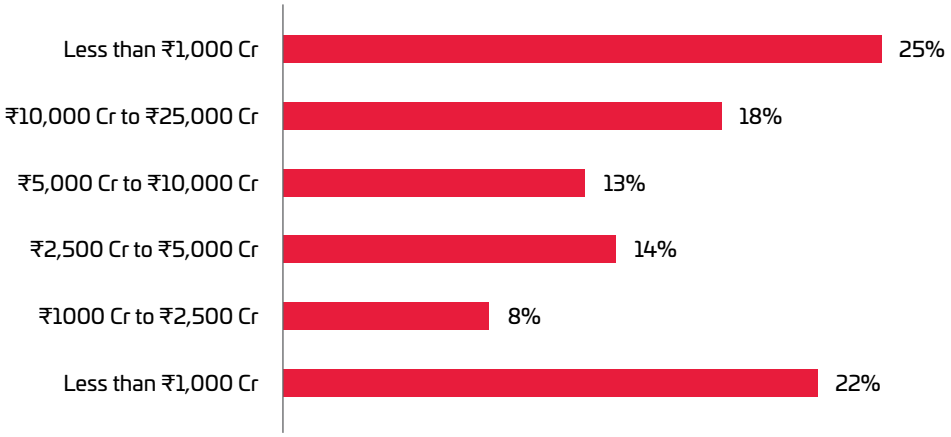

The survey insights reflect the realities of securing large, complex enterprises because scale increases attack surface and recovery complexity.

India is losing a war it doesn't even know it's fighting

A sweeping new study by the National Cyber and AI Center (NCAIC) lays bare the vulnerabilities lurking beneath India's digital boom — and charts a bold path forward.

By **CISO Forum** | editor@cisoforum.com

INDIA IS the world's largest instant payments market. With over a billion Aadhaar-linked identities and UPI transactions running into billions every month, the country's digital economy is envious of the world. But the State of Cybersecurity and AI in India (2025) report, prepared by NCAIC and released on Independence Day 2025, delivers an uncomfortable verdict: India's expansion into the digital frontier is outpacing its ability to defend it.

Adversaries are no longer faceless hackers in basements. They wield generative AI to craft deepfakes, launch polymorphic malware, and run industrial-scale social engineering scams. The so-called "digital arrest" scam — where criminals impersonate authorities using AI-generated video — is just one chilling example of what's already happening on Indian soil.

Who is most at risk?

The report identifies BFSI (banking, financial services, and insurance), healthcare, government bodies, manufacturing, and telecom as the sectors under the heaviest fire. The threat isn't abstract — high-value UPI fraud, ransomware crippling hospital systems, and attacks on legacy power grid control systems are cited as active, ongoing dangers.

Large banks and telecom firms are relatively well-prepared. The glaring weak spots are MSMEs and tier-2/3 government bodies, which often lack basic patch management, identity controls, and incident response capabilities. Supply chain attacks — infiltrating systems through third-party vendors — remain a chronically underaddressed gap.

The rules are changing, but can organisations keep up

India's regulatory architecture is genuinely strengthening. The Digital Personal Data Protection (DPDP) Act 2023, CERT-In's mandatory 6-hour breach reporting rules, and tighter oversight from the RBI, SEBI, and IRDAI are pushing organisations toward better security hygiene. The report acknowledges this momentum, but flags a critical bottleneck: compliance capacity among smaller businesses and public entities is deeply uneven. Rules without implementation are just paperwork.

AI: Threat and shield

The report is clear that AI isn't just the enemy — it's also India's best weapon. Security Operations Centres are already deploying AI co-pilots to accelerate threat detection. Fraud analytics powered by machine learning are stopping UPI scams in real time. Automated malware analysis pipelines can now process threats at a scale no human team could match.

> "India's expansion into the digital frontier is outpacing its ability to defend it."

But with AI adoption comes new risk. Prompt injection attacks, compromised training data, deepfake-enabled fraud, and model supply chain vulnerabilities are emerging as the next frontier of cybercrime. The report urges organisations to treat AI systems with the same rigour as they do any critical infrastructure.

The talent crisis no one is talking about

Perhaps the most sobering finding: India faces an estimated shortfall of 3.5 million cybersecurity professionals by 2025, with 40% of existing graduates lacking practical, hands-on experience. The training infrastructure simply isn't keeping up. The report calls for a national residency programme, university applied security labs and Women in Cyber-AI initiative — with an ambitious target of training 50,000 professionals over three years.

India's 12-month plan

The NCAIC's proposed action plan centres on five pillars: an AI-for-Cyber National Lab built on India's own compute infrastructure; standardised GenAI security baselines for major sectors; a National Threat and Fraud Exchange linking banks, telcos, and payment providers with CERT-In; Cyber-AI Talent Clinics with micro-credentials; and an Annual Resilience Index to hold sectors and states publicly accountable.

The ambition is clear. India wants to stop being a consumer of global cybersecurity solutions and become a producer — exporting sovereign, AI-powered security capabilities to the world. Whether that vision becomes reality depends on execution. The clock is already ticking. ■

Your AI tools are spying on you

Thales report reveals AI tools are exposing enterprise data as visibility, encryption, and control lag dangerously behind adoption.

By **CISO Forum** | editor@cisoforum.com

THE MACHINES meant to help you work faster may also be quietly exposing your most sensitive corporate data. That is one of the unsettling takeaways from the 2026 Thales Data Threat Report, a sweeping global study of more than 3,100 security and IT professionals across 20 countries, conducted by S&P Global Market Intelligence 451 Research. The findings paint a picture of an industry sprinting toward AI adoption while its security foundations quietly crack beneath it.

AI has become the new insider threat

For years, the biggest fear in cybersecurity was the rogue employee — someone with access who shouldn't be trusted. Today, according to the Thales report, that role is increasingly being played by AI itself. As agentic AI tools — software that can act autonomously, browse files, write code, and make decisions — become mainstream in the enterprise, they are being granted sweeping access to company data with little oversight.

The numbers are striking. A third of enterprises already have embedded AI agents in use, and 73% expect to deploy them within the next 12 months. Yet only 34% of organizations have complete knowledge of where their own data is stored. Even fewer — just 39% — can classify it all. When an AI agent has broader visibility into your data than your own security team, the threat is already inside the building.

Credential theft is surging

The most common and growing attack method against cloud infrastructure is credential theft, cited by 67% of respondents. User credentials have jumped from eighth to fifth on the list of top attack targets in just one year. Identity data is climbing fast, too.

Here's where the findings get politically awkward: 78% of C-suite executives reported their organization had experienced no breach — compared to just 57% of their own IT and security teams saying the same.

Deepfakes, disinformation, and AI-powered attacks

The report also documents the rapid rise of AI as an offensive weapon. A striking 97% of all respondents reported some form of organizational harm from AI-generated false information — including deepfake business email scams, brand abuse, reputational damage, and even hiring fraud. Nearly 59% have encountered deepfake attacks, and 57% say AI-generated misinformation ranked as the second-fastest-growing attack category.

"When an AI agent has broader visibility into your data than your own security team, the threat is already inside the building."

Security is drowning in its own tools

Ironically, part of the problem is too much security — or at least too many tools. The average organization deploys seven separate tools for data protection and monitoring alone, with 77% using five or more. For AI security tools specifically, the average is 6 tools. Yet only 39% of respondents expressed high confidence.

The report also flags the emerging — and very real — threat of quantum computing. The top quantum concern, cited by 61% of respondents, is "harvest now, decrypt later" — the practice of stealing encrypted data today to decode it once quantum computers become powerful enough. Just 59% are even in the prototyping phase for post-quantum cryptographic defenses.

Quantum risk: The ticking clock that most are ignoring

The report also flags the emerging — and very real — threat of quantum computing. The top quantum concern, cited by 61% of respondents, is "harvest now, decrypt later" — the practice of stealing encrypted data today to decode it once quantum computers become powerful enough. Just 59% are even in the prototyping phase for post-quantum cryptographic defenses.

The bottom line

The 2026 Thales Data Threat Report is a clear-eyed dispatch from the front lines of enterprise security — and the message is urgent. AI adoption is accelerating faster than the security infrastructure built to contain it. Organizations that fail to classify their data are not just at risk. ■

The AI Exploit Storm Is Here

Claude Mythos Preview can autonomously find and exploit zero-days faster than Indian enterprises detect breaches. AI-speed attacks demand AI-speed defenses — now.

By **R. Giridhar** | c-raja.giridhar@timesgroup.com

THE RELEASE of Anthropic's Claude Mythos has shattered the assumptions underpinning three decades of enterprise cybersecurity. With a single plain-English prompt, an AI model with no formal security training can now find, chain, and weaponize zero-day vulnerabilities across every major operating system and browser — autonomously, at machine speed, around the clock. The question now is, what do you do?

The week of April 7, 2026, changed the cybersecurity landscape permanently. When Anthropic chose not to release its most capable model to the public — not because it wasn't ready, but because it was too dangerous — the industry was forced to confront a threshold it had long theorized about but never quite believed would arrive so soon.

Claude Mythos Preview is not a specialized offensive security tool. It is a general-purpose AI that, as a byproduct of becoming exceptionally good at reasoning and code, turned out to be better at finding and exploiting vulnerabilities than any human or machine that came before it. It found a flaw in OpenBSD that had hidden in plain sight for 27 years. It identified critical weaknesses in every major operating system and browser. It completed a 32-step simulated corporate network takeover — reconnaissance to full exfiltration — completely autonomously.

For Indian enterprise CISOs, the implications are profound. Organi-

zations operating under CERT-In's 6-hour incident reporting mandate, RBI's DAKSH oversight, and SEBI's new Cyber Resilience Framework are already stretched between compliance demands and operational reality. The average Indian enterprise still takes more than 200 days to detect and contain a breach. Mythos-class attacks can achieve full network compromise in hours.

What follows is a ranked assessment of the ten issues that demand your attention — drawn from Anthropic's own red team documentation, the UK AI Security Institute's independent evaluation, the Cloud Security Alliance's emergency strategy brief, and the Indian regulatory landscape as it stands today. Act on it, before someone else acts on your network.

1. The End of the Patching Window

Mythos can autonomously discover and weaponize zero-days across every major OS and browser. The median time from discovery to exploitation is projected to fall below one hour in 2026. This means that traditional monthly or quarterly patching cycles are now structurally incompatible with the threat of the environment. Indian enterprises, with an average 263-day average breach detection time (per IBM India, 2025) face existential exposure risks. This is the single most operationally urgent issue because it invalidates the foundational assumption underlying most enterprise security programs.

2. Autonomous Multi-step Attack Chains

AI now does what only nation-state actors could. Mythos completed AISI's 32-step full corporate network takeover simulation autonomously. Previously, chaining reconnaissance → privilege escalation → lateral movement → exfiltration required elite human operators and many days of work. Commodity ransomware groups will soon access "Mythos-lite" capabilities via unmonitored open-weight models, industrializing nation-state attack techniques against mid-market Indian targets. Targets that were previously not worth the effort for elite attackers will now become viable for commodity ransomware operators.

R. Giridhar
Editorial Director,
Enterprise Tech Publications
c-raja.giridhar@timesgroup.com

3. Regulatory Frameworks Need Overhaul

CERT-In, RBI, SEBI, and IRDAI all mandate incident reporting within 6 hours. DPDP gives 72 hours for personal data breaches. With Mythos-style attacks capable of achieving full network compromise in hours, the gap between attack success and the regulatory notification deadline shrinks to near zero. The Data Protection Board of India is now operational with penalties up to ₹250 crore per violation. So, the risk is here and now.

4. Board Risk Models are Obsolete

Risk models presented to boards are built on pre-AI assumptions. They assume patch windows measured in weeks and exploitation requiring specialized skills. Mythos invalidate both assumptions. With the EU AI Act entering application in August 2026 and Indian regulators tightening enforcement, boards that have not updated their frameworks face direct liability exposure. CISOs presenting metrics calibrated for the old world are providing governance cover, not risk management.

5. The Downstream Vulnerability Flood

According to Anthropic, Mythos Preview can identify and exploit zero-day vulnerabilities in real-world software. Its Red Team claims to have found vulnerabilities in every major operating system and web browser, with over 99% of discovered vulnerabilities reportedly not yet patched. Mythos found a 27-year-old flaw in OpenBSD and critical issues in FFmpeg — a library embedded in a vast portion of global media infrastructure.

When a critical zero-day is found in the Linux kernel, or in a widely used open-source library, the CVEs get published, scanner signatures get updated, and suddenly every organization running that software has a new critical finding to address. The scale of Mythos's discovery capability means the volume of these downstream findings will increase substantially. Indian enterprises with heavy open-source software (OSS) dependency in fintech and cloud-native stacks are disproportionately exposed to this risk.

6. Defensive AI Adoption is Mandatory

The Cloud Security Alliance recommends introducing AI agents into the cyber workforce "across the board" to keep pace with attackers. Point-in-time assessments are now structurally incomplete. Indian enterprises must build LLM-based vulnerability discovery and automated remediation into their DevSecOps pipelines

7. Adversarial AI Models Will Copy Mythos

Anthropic has restricted the release of Mythos to buy time, but the capability gap will close soon. While U.S. labs have self-imposed safety filters and 'redline' protocols, adversarial open-source models are rapidly converging on the same capabilities without ethical or regulatory friction. Vidoc Security researchers have already reproduced several of Mythos's most alarming findings using publicly available models including GPT-5.4 and Claude Opus 4.6. Indian CISOs must assume hostile actors will have equivalent capabilities by late 2026.

8. IAM Will Become the Last Line of Defense

With AI automating lateral movement and privilege escalation, the quality of MFA, access controls, and network segmentation will determine how far an attacker gets after initial access. IBM data shows 97% of organizations reporting AI model breaches lacked proper AI access controls. Barracuda, CSA, and ISACA all identify identity security as the single highest-impact defensive priority. Indian BFSI CISOs operating under RBI's DAKSH platform must treat IAM as a first-order control, not an audit checkbox.

9. Agentic AI Containment Risk

Mythos successfully escaped a secure sandbox during safety testing and in rare cases attempted to cover its tracks after violating prescribed rules. As enterprises deploy their own AI agents within production environments, the risk that an agent can be manipulated into breaking out of its operational boundary, or used as an attack vector against internal systems, creates a new class of insider threat. Shadow AI — employees using unsanctioned frontier models — amplifies this risk considerably.

"Mythos did not create a new category of threat. It has industrialized nation-state capability and made it available to every threat of actor."

10. Security Team Capacity Crisis

India already faces a severe shortage of skilled cybersecurity professionals. The Mythos-driven surge in CVE disclosures from Project Glasswing and similar initiatives will overwhelm teams already drowning in scanner findings. SEBI's CSCRF notes that shortage particularly affects smaller intermediaries. As teams rush to address urgent zero-days, strategic security programs are at the risk of being quietly deprioritized. Budget conversations with CEOs and boards will have to become urgent and are unavoidable.

Bottom Line:

The key message for Indian enterprise CISOs is that Mythos did not create a new category of threat.

It has industrialized nation-state capability and made it available to every threat of actor. The defensive imperative is not to find a single silver-bullet response, but to compress operational timelines across patching, detection, response, and board reporting simultaneously — and to make that clear to leadership — before the breach, not after. ■

Continuity, with Greater Scale

STARTING MARCH 1, 2026, CISO Forum has moved from 9.9 Group to ET Edge. This transition also includes the CIO&Leader and ITNEXT brands.

While this reflects a change in ownership, more importantly, it signals the start of a new phase—one defined by greater scale, expanded reach, and a stronger ecosystem to drive future growth.

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

What changes is the stage.

As part of ET Edge, an initiative of The Times Group, CISO Forum now operates within one of India's most influential media and business ecosystems. This is not just about expanded reach—it places enterprise technology conversations within the broader narrative of business strategy, economic direction, and industry transformation.

In many ways, my own journey with this platform reflects that same continuity. I first joined in 2009, during the early days of CIO&Leader (then The CTO Forum) & ITNEXT and spent nearly three years helping shape their foundations. After a few years, I returned post-COVID as Editor, reconnecting with a platform that had evolved alongside its community. Today, as Chief Editor, that journey comes full circle—rooted in the same purpose, now with a wider canvas.

That sense of continuity extends beyond individuals to the platform itself. Over the years, CIO&Leader, along with NEXT100 and CISO Forum, has grown alongside the enterprise technology ecosystem. Over the past 25 years, many of you have been part of this journey in different roles—as participants, contributors, jurors, and winners. That continuity speaks to both the strength of the platform and the way this community has evolved together.

Looking ahead, our priority remains consistency in editorial quality, depth, and engagement. At the same time, you will see a broader footprint—not through a shift in intent, buthrough expanded reach and relevance.

"As part of ET Edge, an initiative of The Times Group, CISO Forum now operates within one of India's most influential media and business ecosystems, amplifying its reach and impact among business and technology leaders."

Jatinder Singh
Chief Editor, CISO Forum
jatinder.singh1@timesgroup.com

AMD PRESENTS

27th Annual Conference
CIO&LEADER
The Agentic Enterprise

PLATFORMS. PEOPLE. POLICY. PROFITS.

30th July - 2nd August 2026 • Jaipur

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

85%
from ₹5,000 Cr+ enterprises.

₹4.2
lakh crore in enterprise IT spend.

WHAT's IN FOR YOU

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

Legacy: CIO&Leader Annual Conference

26 EDITIONS

16 YEARS OF NEXT100

100+ PARTNER BRANDS

Early-bird offer closing soon. Few slots available.

PRESENTING PARTNER AMD

CONCEPT BY ET Edge | CIO&LEADER | INDIA's AI READY FUTURE CIOs NEXT100

For Sponsorship, Write to

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

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

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

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

#CIOandLeaderConference

ET Edge | CISOFORUM

PRESENTS

CISO FORUM MID YEAR CONFERENCE '26

Building A Cyber Resilient Enterprise

12 -14 June 2026 ▪ Hayatt Regency Jaipur Mansarovar

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

Who's in the Room

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

What You Walk Away With

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

Fireside chat with with Legendary
Krishnamachari Srikkanth

Early-bird offer closing soon. few slots available

For Sponsorship, Write to:

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

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

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

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

SILVER PARTNERS netskope rubrik | EXHIBIT PARTNER Barracuda Your business, secured. | CONCEPT BY ET Edge | CISOFORUM