JUNE 2026
SUPPLEMENT OF CIO&LEADER

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

CISOFORUM

Security For Growth And Governance

TRUSTED SPEED, NOT JUST FASTER DETECTION

THE REAL TEST FOR CISOS IN THE AI ERA

Top Row : **Dr. Siraj Rahim,** MRF | **Utkarsh Sawant,** Diageo
Bottom Row : **Sudipta Biswas,** Emami | **Namrata Bhise,** FIS Global | **Pragnesh Mistry,** RPG Enterprises

ET Edge | CIO&Leader

studiotalks

ET EDGE I CIO&LEADER STUDIOTALKS

WHERE TECHNOLOGY MEETS THE SPOTLIGHT !

ET Edge I CIO&Leader presents StudioTalks - a premium platform where India's most influential CIOs and CTOs take center stage. Captured with high production aesthetics, sleek visuals and dynamic backdrops. StudioTalks transforms leadership insights into an engaging cinematic experience and brings India's top CIOs and CTOs into the spotlight.

This exclusive series explores visionary leadership, emerging technologies and strategic transformation - all presented in a format that blends deep insights with the visual polish of a professional studio production.

WHY JOIN STUDIOTALKS

- Engage in powerful conversations that shape the future of enterprise IT
- Share your expertise in a high-impact, TV style format.
- Be featured among India's top technology leaders

Be the voice of transformation. Be part of ET Edge I CIO&Leader StudioTalks.

SECURE YOUR SPOT NOW !

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

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

Follow us: @CIOandLeader

The privacy reckoning

THE DPDP Act has given Indian enterprises a deadline. But compliance alone won't solve a problem that begins long before the CISO is even in the room.

Privacy has always been somebody else's problem. Legal owned the consent of clauses. HR managed the employee data notices. Marketing is worried about cookies. The CISO was busy with firewalls, incident responses, and board-level anxiety about ransomware. Privacy was an adjacent and related issue, but not a core.

That gap is closing. The Digital Personal Data Protection Act (DPDPA) has made privacy an operational reality — breach of notification obligations, consent infrastructure requirements, data fiduciary accountability, and meaningful financial penalties. Since all these sit at the intersection of data, systems, and risk, they are drifting toward the CISO.

But the hard truth is that being handed the privacy brief is not the same as the CISO being empowered to execute it.

In many Indian enterprises, the CISO reports into the CIO or CFO, holds limited authority over product decisions, and is made part of architecture conversations after the key choices have been made. The prescription of "privacy by design" sounds compelling until you realize that design happens upstream — where the CISO is rarely a regular presence. Having authority upstream requires organizational mandate, not just security conviction.

The regulatory environment adds further complexity. With implementation rules still being finalized and the Data Protection Board yet to be fully constituted, many enterprises are in a holding pattern. They are waiting for enforcement clarity before committing significant privacy investments. The CISO who advocates for a consent management platform faces an uphill effort.

But none of these diminishes strategic imperatives. CISOs who are making the privacy shift — from reactive compliance to proactive data governance — are building capabilities that will outlast any specific regulation. But their ability to lead and influence depends on something the CISO cannot create unilaterally: a CEO and board who decides that privacy is a business value, before a regulator insists it must be. ■

"Being handed in the privacy brief is not the same as the CISO being empowered to execute it."

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

CONTENTS

JUNE 2026

COVER STORY

08-15

Trusted Speed, Not Just Faster Detection

Frontier AI compresses threat detection from months to hours. But speed alone won't satisfy a board, an auditor, or India's Data Protection Board. Five senior CISOs explain why trusted speed, not raw detection, will decide who survives the DPDP Act's 72-hour clock, and who doesn't.

Cover Design by:
Manish Kumar

INTERVIEW

INSIGHTS

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**
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**
Senior Producer: **Sunil Kumar**

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

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

Editor: **Vikas Gupta**

9.9 GROUP

Rohit Sharma joined Siemens Energy as Head of Cybersecurity

Rohit Sharma has joined **Siemens Energy India** as **Head of Cybersecurity,** leading initiatives in cyber resilience, risk management, and infrastructure protection. Previously, he held cybersecurity leadership roles at Adani Power, Adani Energy Solutions, and NTPC. With over 20 years of experience, he brings deep expertise in securing critical energy infrastructure.

Ashwini Kumar Choudhary joins RBL Bank as Chief Risk Officer

Ashwini Kumar Choudhary has been appointed **Chief Risk Officer-Designate** at **RBL Bank,** where he will strengthen enterprise risk management and governance. Previously, he served as Group CRO at Union Bank of India and held risk leadership roles at Tata Cleantech Capital, Srei Infrastructure Finance, HSBC, HDFC Bank, and Bank of Baroda.

Hemang Doshi joins Hitech i2i as Director – IT Infrastructure & Security

Hemang Doshi has assumed an additional role as Director – IT Infrastructure & Security at Hitech i2i, leading initiatives in infrastructure, cybersecurity, and resilience. He continues at Hitech Digital Solutions and previously held leadership roles at Apexon, InfoStretch, and other organizations. He brings extensive expertise in cybersecurity, cloud, compliance, and enterprise transformation.

Vijay Pukale joins Vivicta as Lead Security Advisor

Vijay Pukale has been elevated to **Lead Security Advisor** at **Vivicta**, where he will help customers strengthen business-aligned cybersecurity programs and governance. Previously, he led security, cloud, and risk initiatives at Vivicta and held leadership roles at Cloud4C, TietoEVRY, Mphasis, Tech Mahindra, Larsen & Toubro, and ISRO.

Proofpoint launches active exploits protection

New solution helps organizations focus on vulnerabilities attackers are actively exploiting

By **CISO Forum** | editor.tech@timesgroup.com

PROOFPOINT HAS unveiled Active Exploits Protection, a new cybersecurity solution designed to help organizations defend against the growing wave of AI-accelerated cyber threats. The platform identifies actively exploited vulnerabilities and enables organizations to take immediate action before those threats cause damage.

The launch comes as advanced AI systems dramatically shorten the time between vulnerability discovery and real-world exploitation. According to Proofpoint, attacks can now begin within hours—or even minutes—of vulnerabilities being identified, leaving traditional patch-based security approaches struggling to keep pace.

Intelligence based on real-world attacks

The new solution is powered by Proofpoint's extensive threat intelligence network, which analyzes more than two billion emails daily and draws insights from over 5,000 global sensors. In 2026 alone, the company generated more than three million exploit-related alerts and identified 12 actively exploited vulnerabilities, exceeding the number currently tracked in public exploitation catalogs.

"Organizations need to understand what attackers are exploiting in real time and reduce their exposure immediately," said Sumit Dhawan, CEO of Proofpoint.

> "Organizations need to understand what attackers are exploiting in real time and reduce their exposure immediately."
>
> **Sumit Dhawan,** CEO, Proofpoint

Prioritizing what matters most

Proofpoint says fewer than 6% of disclosed vulnerabilities are ever exploited in real-world attacks. Yet, security teams often spend significant time addressing thousands of alerts based on severity scores rather than actual risk.

Active Exploits Protection aims to solve this challenge by helping teams prioritize vulnerabilities that attackers are actively targeting. The platform can automatically convert exploit intelligence into protective measures in about 35 seconds, significantly reducing exposure of windows for zero-day and emerging threats.

The solution also supports AI-driven security workflows, integrates with existing SOC and vulnerability management tools, and enables organizations to focus resources on risks with the greatest business impact. ■

CrowdStrike expands QuiltWorks to tackle AI cyber risk

Cyber insurers join initiative to help organizations manage AI-driven security and financial exposure

By **CISO Forum** | editor.tech@timesgroup.com

CROWDSTRIKE HAS expanded its Project QuiltWorks initiative by bringing together leading cyber insurance and risk management firms, including Coalition, Liberty Mutual Insurance, Lockton, Resilience, and Marsh. The move marks the next phase of the framework, extending its focus from identifying frontier AI security risks to helping organizations manage the financial impact of those risks.

As advanced AI models accelerate vulnerability discovery and shorten the time between identification and exploitation, businesses are facing increased operational and financial exposure. CrowdStrike said Project QuiltWorks is designed to provide a coordinated approach that combines risk discovery, remediation, and financial protection.

"Frontier AI risk doesn't stop at technology; it lands on the balance sheet," said Daniel Bernard, Chief Business Officer at CrowdStrike. He noted that the framework integrates technology, remediation services, and financial safeguards to help organizations address emerging AI-related threats.

> "Frontier AI risk doesn't stop at technology; it lands on the balance sheet."
>
> **Daniel Bernard,** Chief Business Officer at CrowdStrike

A new model for AI risk management

Powered by frontier AI models from OpenAI and Anthropic, Project QuiltWorks combines CrowdStrike's AI-driven vulnerability discovery and threat prioritization with remediation services from systems integrators and the expertise of cyber insurers.

The expanded framework introduces financial risk modeling, underwriting intelligence, and coordinated mitigation capabilities. It aims to help organizations understand their financial exposure, prioritize the most critical risks, strengthen their eligibility for cyber insurance, and improve resilience against AI-driven cyber threats.

Industry partners highlighted the growing challenge posed by frontier AI, emphasizing the need for continuous visibility, proactive risk management, and stronger collaboration between cybersecurity providers and insurers. The initiative seeks to create a continuous feedback loop that connects security intelligence, remediation efforts, and insurance insights to reduce both cyber and financial risk. ■

Palo Alto Networks secures the rise of AI agents

Palo Alto Networks will acquire AI Gateway pioneer Portkey to secure and govern enterprise autonomous AI agents.

By **CISO Forum** | editor.tech@timesgroup.com

PALO ALTO Networks the global cybersecurity leader, has announced its intent to acquire Portkey, a pioneer in AI Gateways. Portkey provides a centralized control plane to manage and protect autonomous AI agents, already processing trillions of tokens monthly with the low latency required for agent-to-agent communication.

As enterprise AI adoption shifts from copilots and applications to autonomous agents, the security gap has widened sharply. These agents function as highly privileged insiders, executing vast volumes of automated decisions across internal and external systems. Post-acquisition, Portkey will serve as the AI Gateway for Prisma AIRS™, acting as a central nervous system to monitor, route, and secure every AI transaction enterprise-wide.

"As autonomous agents join the enterprise workforce, they also become a new, unmanaged attack surface."

Lee Klarich, Chief Product & Technology Officer, Palo Alto Networks

Lee Klarich, Chief Product & Technology Officer at Palo Alto Networks, said that integrating Portkey into Prisma AIRS will enable organizations to deploy and govern AI agents confidently, gain visibility into agentic traffic, and control and protect against agentic threats.

The unified platform is designed to: secure AI interactions through runtime policy enforcement and least-privilege identity controls; ensure mission-critical reliability with 99.99% uptime via semantic routing, automated failovers, and full audit telemetry; and enable global AI governance through centralized artifact management, version control, and cost optimization across 3,000+ LLMs and MCP tools.

Rohit Agarwal, CEO and Co-Founder of Portkey, said the deal balances developer flexibility with security control, positioning the AI Gateway as foundational infrastructure for enterprises deploying autonomous agents with confidence.

Palo Alto Networks will continue supporting existing and new Portkey customers post-close. The transaction, subject to customary closing conditions, is expected to close in Palo Alto Networks' fourth quarter of fiscal 2026. ■

TRUSTED SPEED, NOT JUST FASTER DETECTION

AI's real test isn't how fast an organization spots a threat. It's whether it can decide, act, and later prove it acted correctly — inside a regulatory clock it no longer controls. That kind of speed only survives contact with regulation when governance is structural; not a checklist applied after the fact.

By **Raja Giridhar & Jagrati Rakheja**
c-raja.giridhar@timesgroup.com

"The risk is not that you will be breached. The risk is that you will know you are breachable and fail to act in time to avoid liability. That is a governance risk. Governance failures are director liabilities."

IT WAS the third slide of the presentation when the chairman stopped her.

The CISO of a Mumbai-based BFSI conglomerate had just shown the board a heat map: 14,000 new vulnerabilities discovered across the group's digital estate in the previous quarter, mapped by an AI-augmented scanning tool that had replaced the legacy vulnerability management platform. The number was not shocking. The shock was the annotation in the corner: "Estimated financial exposure if 5% of critical vulnerabilities are exploited within 72 hours of discovery: ₹340 crore."

"Where does that number come from?" the chairman asked. "And why are we only seeing this now?"

The CISO had prepared for this. The financial model had been trained on breach-cost datasets from the Indian market, adjusted for DPDP Act 2023 penalty exposure—up to ₹250 crore for failure to maintain reasonable security safeguards, ₹200 crore for failure to notify the Data Protection Board within 72 hours. The tool had not just found the holes. It had priced them.

But the chairman's second question was the one that mattered. "If we can see this, who else can?"

That question is now echoing through boardrooms across India. Anthropic's Project Glasswing—deploying Claude Mythos Preview to 150 organizations across 15+ countries—is the current benchmark. Whether it delivers on its promise at scale remains to be seen. What is certain is that the capability category it represents is now irreversible. Within 6–12 months, multiple AI vendors will field comparable offensive and defensive capabilities. Some will arrive with safeguards. Others will not. The open-source replication is already underway.

For Indian enterprises, this is not a single-vendor story. It is a structural shift. AI models can now discover vulnerabilities at machine scale, generate patches, simulate attacks, and refactor legacy code. The bottleneck is no longer discovery. It is triage, validation, and remediation capacity—and the compliance of velocity required to act before liability crystallizes.

The 12-month warning is simple: the enterprises that move now will define the defensive standards. The rest will inherit the risk. And the risk is not technical. It is financial, regulatory, and competitive.

The Five Questions Every Indian board must ask its CISO before the next audit cycle.	
1	What is the estimated financial exposure of our unremediated vulnerabilities, and how does it compare to our AI-security investment?
2	If an AI-augmented attacker maps our external attack surface in hours rather than weeks, which control fails first—and have we tested this?
3	What is our compliance velocity under the DPDP Act 2023 and CERT-In directives: can we detect, validate, and report within the mandated windows?
4	What is the minimum viable AI-security stack for our organization size, and what manual processes does it replace in year one?
5	What governance structure ensures AI-security investment survives beyond the pilot phase—and who owns the business case?

The frontier AI landscape

Anthropic's expansion of Project Glasswing is the headline. The subtext is what happens next. Within 6–12 months, at least four other major AI vendors are expected to field comparable capabilities, including two with significant India presence. Some will arrive with the safeguards Anthropic has built into Mythos—restricted model access, audit logging, human-in-the-loop requirements. Others will not.

The open-source replication is already underway. Independent research groups have demonstrated models with 70% of Mythos-class performance using publicly available datasets, and the timeline for ungoverned deployment is measured in months, not years. That democratization cuts both ways. Attackers will acquire these capabilities faster than enterprises, with fewer constraints and no compliance overhead.

For Indian enterprises, the vendor-selection dilemma has strategic consequences. The organization that adopts a well-governed AI security platform in Q3 FY26 will build institutional knowledge, tune detection models to local threat patterns, and integrate with existing SOC workflows. The organization that waits for a perfect standard will face the same threat of landscape with the same legacy tools—and a smaller talent pool to operate them.

The India-specific maturity gap is stark. Tier, I bank and large IT/ITES firms have piloted AI-augmented vulnerability management since late 2024. Mid-market manufacturers, PSUs, and critical infrastructure operators are still evaluating whether their on-premises and OT environments can accommodate AI security tools at all.

The changed risk calculus

Here's the equation every CISO learned in certification school: probability × impact = residual risk. It is now wrong.

Not because math changed. Because the variables changed non-linearly. A human-led vulnerability assessment might identify 200–400 issues in a mid-sized enterprise estate over a quarter. An AI-augmented scan can surface 10,000–15,000 in the same period, with severity scoring that correlates

Sudipta Biswas
Chief Information Security Officer, Emami

"I would frame AI-security maturity to the board as a business trust and resilience capability, not as a cybersecurity technology project. Boards do not need technical fear. They need business clarity. The language should not be, 'AI attacks are coming and we are at risk.' The more effective board-level message is: 'AI is changing the speed of cyber risk. Our competitive advantage will depend on whether we can make trusted decisions faster than the threat can turn into business loss.' A simple way to explain this to the board is: "Cybersecurity today is like air-traffic control. We have many signals coming from endpoints, email, cloud, identity, SaaS, network, applications, and third parties. Some are routine, some are duplicates, some are false alarms, and some are real threats. AI can help us read the radar faster, but faster radar alone will not make us safer unless we can also decide and act faster."

exploitability to business-critical asset exposure. The bottleneck is no longer discovery. It is triage, validation, and remediation capacity.

Boards don't need to understand the cost structure. They need to understand that the cost structure has changed. Every unremediated critical vulnerability is a contingent of liability. Under the DPDP Act 2023, that liability has a statutory price. The penalties are severe: up to ₹250 crore for failure to maintain reasonable security safeguards; up to ₹200 crore for failure to notify the Data Protection Board within 72 hours of becoming aware of a personal data breach. The Act's phased implementation—core provisions effective 13 May 2027—means boards have a compliance runway, but also a countdown.

The risk is not that you will be breached. The risk is that you will know you are breachable and fail to act in time to avoid liability. That is a governance risk. Governance failures are director liabilities. The CISO's job is to ensure directors have the information to avoid them.

The cost of AI security

The investment case for AI security tooling is a capital allocation decision with competing claims on limited budget. CISOs will need to make this case without assuming a net budget increase, because for most organizations, there is none.

The matrix is simple: spend here, avoid cost there. For a large enterprise with ₹5,000 crore annual revenue, a single DPDP Act penalty at the upper threshold (₹250 crore) represents 5% of revenue. The cost of an AI-augmented vulnerability management platform, including infrastructure upgrade and workforce transition, typically falls in the range of ₹8–15 crore annually for a comparable estate. Math is not subtle.

But the hidden costs are where boards are surprised. Infrastructure upgrade for AI security tools is not a software purchase. It is a compute, storage, and network architecture decision. AI-scale scanning generates data volumes that strain the legacy of SIEM deployments. On-premises estates require GPU or TPU capacity that most Indian enterprises have not provisioned. Hybrid and multi-cloud environments need API integration and data sovereignty compliance that add 20–30% to implementation timelines.

Workforce transition costs are equally underestimated. AI security tools do not eliminate SOC analysts. They redeploy them from detection to validation, from alert triage to threat hunting, from manual patching to automated remediation oversight. Retraining ROI typically turns positive in the second year, though few organizations measure it. The transition period requires parallel staffing that mid-market organizations cannot afford.

For mid-market CISOs, cost modeling is more constrained. Annual cybersecurity spend as a percentage of IT budget in Indian mid-market firms averages 6–8%, against 12–15% in large enterprises. The minimum viable AI-security stack—a combination of open-source AI-augmented scanning, managed detection and response (MDR) with AI overlay, and cloud-native SIEM—can be deployed for ₹80 lakh–₹2 crore annually. But it replaces only a subset of manual processes in year one, and the break-even point depends on whether the organization avoids even one material breach.

Dr. Siraj Rahim
DGM-Operation Technology & Information Security, MRF

"We are entering a phase where AI is redefining not just the scale of cyber risk, but the speed of consequence. In India, frameworks like CERT-In's 6-hour reporting and the DPDP Act's 72-hour mandate signal a clear shift—from reactive cybersecurity to continuous, accountable readiness. The question is no longer whether you can detect a breach, but whether you can act decisively and explain that action with confidence. AI compresses the time to exploit; regulation compresses the time to accountability. The board's mandate is to ensure trusted speed—where the organisation can act fast and prove it acted correctly."

The boardroom playbook

The CISO's job in the boardroom is not to explain how AI finds vulnerabilities. It is to explain why the organization's risk of posture has changed in ways that require immediate capital allocation, governance restructuring,

What Does it Cost?

AI-security investments and constraints vary by industry.

BFSI (Tier I/II Banks): RBI mandates 24×7 SOC, half-yearly VAPT, 6-hour CIMS reporting. AI tools compress VAPT from months to weeks, but RBI does not yet recognize AI-generated reports as audit-compliant. Parallel human-validated assessments with double short-term cost. Typical AI-security investment: ₹10–18 crore/year.

IT/ITES: Client trust is a revenue driver. AI-security maturity is a competitive differentiator in RFPs. Typical investment: ₹6–12 crore/year. Hidden cost: client-mandated penetration testing that duplicates AI findings.

Manufacturing & Critical Infrastructure: OT environments with 20-year asset lifecycles require network segmentation and proxy architectures before AI tools can deploy. Additional infrastructure cost: ₹3–5 crore. Typical total investment: ₹8–15 crore/year.

PSUs: Government procurement rules privilege lowest-cost compliant bids. Premium AI tools struggle against commodity alternatives that meet baseline CERT-In requirements but lack predictive capability. Typical investment: ₹4–8 crore/year, often via phased tender.

or competitive repositioning. Here is the language for three conversations that matter.

- **To the chairman: Risk translation** Boards understand financial exposure, not CVSS scores. A critical vulnerability in a customer-facing payment gateway is not a '9.8 CVSS. It is a '₹45 crore contingent liability with a 72-hour regulatory notification obligation to the Data Protection Board. The language changes the conversation from IT operations to risk management.
- **To the CFO: Investment justification** The CFO does not care about MTTD. The CFO cares about avoiding costs. Frame this way: AI-augmented threat detection reduces mean time to detect from 197 days to 24 hours. At India's average breach cost of ₹18.4 crore, every month of detection delay costs ₹1.5 crore in containment and response. The platform pays itself if it prevents one material breach in 36 months. We have had three near-material incidents in the past 24 months.'

One IT/ITES CISO framed the ask to his CFO as: 'This is not a technology purchase. It is an insurance premium with a known payout structure. The alternative is self-insurance against a liability we cannot quantify.'

- **To the audit committee: compliance velocity** The audit committee understands deadlines. What they often miss is velocity—the speed at which the organization can detect, validate, report, and remediate. The CERT-In 6-hour breach of reporting requirement is not an abstract obligation. It is an operational reality that AI security tools can either accelerate, or complicate, depending on whether the organization's logging, detection, and reporting infrastructure was designed for machine-scale output. A SOC that generates 10,000 alerts per day cannot manually report within 6 hours.

One manufacturing CISO addressed her audit committee: Compliance velocity is now a function of automation maturity. We are currently at 40% automation. The regulator expects 6-hour reporting. The gap is a governance risk, not a technical one.'

Utkarsh Sawant
Global Head of Cyber Strategy & Risk, Diageo

"We translate technical vulnerabilities into business risk by estimating potential financial impact across operational disruption, regulatory exposure, incident response costs, and reputational damage. For board discussions, vulnerability counts are less effective than metrics such as 'risk reduction achieved,' 'time-to-remediation,' and 'potential business loss avoided.' The conversation resonates most when framed in terms of business resilience, regulatory compliance, and financial exposure rather than technical severity alone."

Governance structures that sustain investment

Three models have proven effective for sustaining AI-security investment

beyond pilot phase, scaled to organization size:

- **Board cyber committee with quarterly AI-risk review** (large enterprise, 5,000+ employees): DPDP Act penalty exposure as standing agenda item; CISO reports directly to committee, not through CIO.
- **CISO-CFO collaboration model with joint ownership of security ROI metrics** (mid-market, 500–5,000 employees): Shared dashboard of avoided cost, MTTD/MTTR trends, and regulatory compliance velocity. If there is no joint ownership, who owns the business case for AI security to spend?
- **Third-party risk framework with AI-security maturity scoring** (all sizes): Vendor contracts include SBOM review, AI-assisted security validation, and breach of liability clauses.

Regulatory anticipation

Boards must be briefed on compliance with velocity, not just deadlines. The DPDP Act 2023's core provisions take effect on 13 May 2027. CERT-In's 2025 Cyber Security Audit Guidelines mandate annual audits with CVSS + EPSS scoring, red teaming, and SBOM review. RBI's cybersecurity framework is expected to incorporate AI-assisted security tool guidance by Q4 FY26. The CISO who flags these obligations before compliance deadlines positions security as proactive governance, not reactive compliance.

CISO takeaways

The 12-month warning is not uniform instruction. It is a differentiated imperative. Here is where to start, by organization size.

For Large Enterprises (5,000+ employees)

- Commission for an AI-security readiness assessment by Q2 FY26. Use CERT-In's 2025 audit guidelines as the baseline, not as a ceiling.
- Pilot AI-augmented vulnerability management on a non-critical estate to build institutional knowledge before regulatory mandates require it.
- Establish board cyber committee quarterly review of AI-risk metrics, with DPDP Act 2023 penalty exposure as a standing agenda item.
- Negotiate AI-security maturity clauses into third-party vendor contracts, with SBOM review requirements.

For mid-market organizations (500–5,000 employees)

- Define the minimum viable AI-security stack by Q1 FY26: open-source AI-augmented scanning + MDR with AI overlay + cloud-native SIEM.
- Identify the three manual processes that consume the most SOC analyst hours and evaluate AI replacement feasibility.
- Engage a managed security service provider (MSSP) with AI capability rather than building in-house, given talent constraints.
- Map DPDP Act 2023 obligations to current data handling practices and identify notification workflow gaps before the 13 May 2027 deadline.

For all organizations

- Test your existing controls against AI-augmented adversary speed. If an attacker can map your external attack surface in hours, which control fails first?
- Build a skeptic diagnostic into your AI-security procurement.
- Frame AI security investment as competitive positioning, not cost center management. The board that sees security as a revenue protector allocates differently than the board that sees it as a compliance cost.

Namrata Bhise
Director - Cybersecurity, FIS Globals

"Balancing compliance speed with legal liability in the age of frontier AI and accelerated regulatory mandates—like the CERT-In 6-hour reporting and DPDP Act 72-hour notification—requires shifting from a reactive 'checkbox' mentality to an embedded, automated governance model. To meet tight compliance windows without increasing liability, compliance must be embedded into the technology lifecycle. The proliferation of AI tools has also changed the hiring narrative: it allows us to hire bright minds in tier-2 cities who can be upskilled into 'AI Security Analysts' rather than traditional SOC analysts."

Looking ahead

The frontier will not stay frontier for long. AI-augmented vulnerability management will be standard tooling across Indian enterprises by 2028 — the real question is who is in the room when the standard gets written. Three trajectories will decide that.

The Skeptic's Diagnostic

A 10-question self-assessment before you sign that multi-year AI-security contract.

1	Can your SOC absorb a 10× increase in findings without alert fatigue?
2	Do you have validated false-positive rates for your estate, or vendor benchmarks only?
3	Has the model been tested against adversarial inputs? What was the miss rate?
4	Can your remediation pipeline scale to match discovery velocity?
5	What is your current alert-to-investigation ratio, and how will the AI tool change it at 30, 90, and 180 days?
6	Does the tool provide explainable output analysts can validate, or is it a black box?
7	Have you measured the operational cost of a false positive in analyst hours?
8	Can you export vulnerability data, threat intel, and remediation history in a standard format if you switch to vendors?
9	Does the pricing model penalize scale (per-asset, per-scan, per-alert) as your estate grows?
10	If the vendor is acquired or discontinuing the product, what is your transition plan and estimated cost?

Scoring: count 'Yes' or 'Has Plan' responses.

- 8–10: Proceed with confidence.
- 5–7: Proceed with conditions—pilot terms, explainability demands, contract caps.
- 0–4: Pause and address gaps before committing

General access. AI security capabilities are getting cheaper as open-source models close the gap with frontier ones; independent groups already claim roughly 70% of Mythos-class performance, with broader availability expected by late 2026. For India's mid-market, the constraint will stop being cost and start being capacity can the SOC actually act on machine-scale intelligence, or does it just drown in a faster stream of alerts?

Regulatory maturation. By 2027, the DPDP Act's core provisions will be in force, CERT-In's 2025 audit guidelines will have run a full compliance cycle, and the RBI is expected to have folded AI-assisted security tooling into its supervisory framework. None of these standards exist in finished form yet. The enterprises piloting AI security today are, whether they intend it or not. Supplying the case studies, regulators will lean on when they finish writing the rules.

Defender-attacker equilibrium. The least settled question is who the arms race actually favors. AI tools could give defenders a durable edge — or attackers, unconstrained by procurement cycles, audits, or board sign-off, could simply out-

Pragnesh Mistry
Head – IT & Cybersecurity,
RPG Enterprises

"Effective governance models include board-level cyber oversight, strong CISO-CFO collaboration for risk-based budgeting, cross-functional AI governance committees, and structured third-party risk management using frameworks such as NIST AI RMF and ISO 27001. These approaches help sustain AI-security investment by aligning security initiatives with business objectives, risk reduction, and compliance requirements."

iterate them. CERT-In's AI-driven situational awareness systems, cited in the World Economic Forum's Global Cybersecurity Outlook 2025, point toward one answer: collective defense, with enterprises sharing threat intelligence in something closer to real time. That only works once participation stops being voluntary—and for most sectors, it still is.

Back in that Mumbai boardroom, the chairman's real question — "If we can see this, who else can?" — was never about the scanning tool. It was about whether the organization could move at the speed of its own visibility now demanded of it. That is the test every Indian board will face over the next twelve months: not whether they can see the risk, which AI has made almost unavoidable, but whether they can act on it before the law, the market, or an adversary forces the question for them. ■

Trend	Focus	Target benchmark
Capability Diffusion	Multiple vendors will field Mythos-class models within 6–12 months; some without safeguards	Urgency for budget allocation
The Asymmetric Threat Reality	Attackers will acquire comparable capabilities, likely faster and with fewer constraints	9Need for preemptive defensive investment
The Legacy Infrastructure Trap	AI security tools expose gaps in on-premises, OT, and hybrid environments	Capital expenditure and modernization required
The Workforce & Operating Model Shift	SOC analyst redeployment, new AI-ops functions, retraining ROI; limited AI-security skills	Human capital and organizational design decisions
The Regulatory & Liability Frontier	Emerging standards for AI-assisted security, data sovereignty, breach disclosure obligations under DPDP Act 2023 and CERT-In directives	Compliance risk and director liability exposure
The Supply Chain Cascade	Third-party and open-source risk amplification when vendors lack AI-security maturity	Procurement policy and vendor risk management
The Skeptic's View	AI-generated false positives at scale, alert fatigue, degraded human analytical skills, vendor lock-in to AI security platforms	Balanced risk assessment and prudent investment framing

From tools to trust: Why platformization is now India's cybersecurity imperative

Swapna Bapat, Vice President and Managing Director, India and SAARC, Palo Alto Networks makes the case for platformization, agentic SOC governance, and treating cybersecurity as a board-level business resilience function.

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

INDIAN ENTERPRISES are caught in a familiar trap — managing sprawling arsenals of 40 to 60 disconnected security tools while facing threats that move faster than any analyst can keep up. As boardrooms shift from asking how many controls are deployed to how effectively risk is being managed, the pressure on security leaders to simplify, consolidate, and govern intelligently has never been greater.

Swapna Bapat, Vice President and Managing Director, India and SAARC, Palo Alto Networks, sits at the intersection of that transformation. In this wide-ranging conversation with CISO Forum, she addresses the hard questions — from realistic entry points for platformization and agentic SOC governance to DPDP compliance gaps and the underestimated dangers of AI-driven supply chain exposure. Her message to Indian boards is clear: cybersecurity is no longer a technology line item. It is the foundation of business resilience.

CISO Forum: For Indian enterprise managing 40–60-point solutions, what is the realistic entry point for platformization — and how does PANW quantify the risk reduction to a board still approving budgets line-by-line?

SWAPNA BAPAT: Most Indian enterprises don't begin platformization with a large-scale rip-and-replace exercise. The realistic starting point is usually consolidating areas where operational fragmentation creates the most risk — typically SOC operations, cloud security, identity, or network visibility.

Conversation with boards is also changing. Security leaders are increasingly moving away from positioning cybersecurity as isolated technologies spend and instead framing it as a business resilience and operational efficiency issue. When organizations manage 40–60 disconnected tools, the hidden costs often manifest slower response times, duplicated workflows, alert fatigue, and limited visibility across environments.

"The biggest challenge enterprises are facing today is that AI adoption is moving significantly faster than governance maturity."

Platformization helps reduce operational complexity. The outcomes boards tend to focus on are measurable improvements such as faster detection and response, shorter incident investigation time, reduced tool overlap, and greater visibility across the attack surface. Ultimately, the discussion shifts from "how many tools are we buying?" to "how effectively are we reducing risk and improving resilience?"

CISO Forum: What are the three most operationally dangerous AI risks Prisma AIRS is detecting in Indian enterprise deployments today — and how do you govern AI risks the security team didn't build?

SWAPNA BAPAT: The biggest challenge enterprises are facing today is that AI adoption is moving significantly faster than governance maturity. Business teams are deploying AI capabilities directly into workflows to improve productivity and automate, often before security teams have full visibility into how these systems interact with enterprise data and applications.

A critical first step, therefore, is discovery. Organizations need visibility into which AI applications, models, and tools are already in use across the enterprise, so they can begin sanctioning usage, enforcing governance controls, managing access, and reducing risks of sensitive data exposure and loss.

One of the clearest risks emerging today is uncontrolled access to sensitive information through AI-connected environments and cloud services. Organizations are increasingly discovering that machine identities and non-human access paths are expanding far faster than traditional governance models were designed to handle.

The second major concern is the granting of excessive permissions to AI agents and connected systems. At the same time, the third is the rapid growth of unsanctioned or shadowed AI usage across business functions. These risks become difficult to govern because security teams are often not the owners or builders of these deployments.

CISO Forum: How should a CISO sequence the move to an autonomous SOC, and what governance must be in place before AgentiX takes remediation actions within BFSI or manufacturing infrastructure?

SWAPNA BAPAT: TThe transition toward an autonomous SOC should proceed gradually rather than through immediate full-scale automation. Most organizations begin by introducing AI into investigative workflows, alert prioritization, and threat of correlation before gradually enabling automated response actions in tightly governed environments.

In sectors like BFSI and manufacturing, governance becomes absolutely critical before AI systems are allowed to take remediation actions. CISOs need clear escalation paths, human approval workflows for high-impact decisions, strict role-based access controls, and complete audit visibility into every automated action taken inside the operational environment.

This is becoming especially important because identity weaknesses continue to play a major role in enterprise incidents globally. Unmanaged service accounts, excessive permissions, and fragmented identity governance allow attackers to move laterally much faster once initial access is gained.

CISO Forum: What specifically does Cortex XSIAM do in the first 60 minutes of a ransomware attack that a traditional SIEM-plus-manual-SOC cannot?

SWAPNA BAPAT: In the first 60 minutes of a ransomware attack, speed and correlation matter more than the volume of alerts. Traditional SIEM-led environments often depend on analysts manually stitching together telemetry from multiple tools before action can be taken. That creates delays at exactly the moment attackers are trying to escalate privileges, move laterally, turn off controls, and encrypt assets.

What platforms like Cortex XSIAM are designed to do differently is automate large parts of that early-stage detection and response cycle. Instead of treating alerts as isolated events, the platform correlates activity across endpoints, network traffic, and identity, the cloud, and user behavior in near real time to quickly build an attack context.

In practical terms, that means identifying suspicious lateral movement, detecting privilege escalation patterns, isolating compromised endpoints, prioritizing high-confidence incidents, and automating containment workflows before the attack fully propagates.

"Cybersecurity is no longer a technology line item. It is the foundation of business resilience."

The value is not just in achieving full visibility, but in compressing the time between detection, investigation, and response. In a ransomware incident, those minutes can determine whether the threat is contained early or escalated into a business-wide disruption.

CISO Forum: What is the migration risk that boards don't ask about but should — and how long does a realistic XSIAM migration take for a 10,000-seat Indian enterprise?

SWAPNA BAPAT: The migration risk boards often underestimate is not the technology of migration itself — it is the operational and process of migration around it. Most large enterprises have spent years building analyst workflows, alerting logic, escalation paths, compliance reporting, and integrations around their existing SIEM environments. If that transition is poorly handled, organizations can end up with visibility gaps, alert fatigue, or duplicated operations during the overlap period.

Another challenge is data quality. Many enterprises discover during migration that they are ingesting enormous amounts of low-value telemetry that adds cost and complexity without improving detection outcomes. Migration becomes an opportunity to rationalize what actually matters from a security and business-risk perspective.

The timeline for a full migration depends on factors such as the environment's complexity, the number of integrations, regulatory requirements, and the SOC's maturity. Organizations that already have strong visibility into their workflows, cleaner telemetry, and well-defined operational processes typically experience a much smoother, faster transition.

CISO Forum: How does PANW's India cloud investment concretely change the compliance architecture for BFSI and healthcare — and what DPDP decisions are Indian CISOs still deferring?

SWAPNA BAPAT: The India cloud investment changes the conversation from "can we use cloud-delivered security?" to "how do we use it while keeping data location, latency, and regulatory expectations in view?" For sectors like BFSI and healthcare, that matters because they deal with highly sensitive customer, financial, and patient data, and are under greater scrutiny on where data is processed, logged, and analyzed.

Palo Alto Networks has expanded India-region support across areas such as WildFire, Prisma Access Cloud Management, and Prisma AIRS runtime/API detection services, which help customers address data residency and performance requirements more directly.

This investment has also expanded across a broader set of cloud-delivered security capabilities, including WildFire, Advanced WildFire, Prisma Access, Autonomous Digital Experi-

ence Management (ADEM), Cloud Identity Engine, Advanced URL Filtering, DNS Security, SaaS Security, Prisma Cloud, Cortex XSIAM, and Prisma AIRS runtime and API detection services. This helps organizations address data residency, latency, operational visibility, and performance requirements more effectively within India.

That said, DPDP readiness cannot be achieved solely through local cloud infrastructure. Many CISOs are still defining how to manage sensitive data, telemetry, retention, access governance, and breach of reporting obligations across the enterprise.

CISO Forum: How does single-vendor SASE change the security operations model — and where does it break down in practice for large Indian conglomerates?

SWAPNA BAPAT: Single-vendor SASE changes the operating model by bringing access, policy, and visibility into one more consistent layer across users, branches, applications, and cloud environments. For security teams, that means fewer handoffs between networking and security teams, fewer policy gaps, and faster investigation when something goes wrong.

The value is especially clear for large, distributed enterprises, where users are no longer sitting behind a single corporate perimeter. The security model has to follow the user, device, application, and data wherever they are.

Where it breaks down in practice is not usually the technology promise, but the operating reality. Large Indian conglomerates have multiple business units, legacy infrastructure, inherited vendors, and different compliance needs. The challenge is getting enough organizational alignment to standardize policies, reduce duplication, and implement SASE consistently across very different parts of the business.

> "The migration that often underestimate is not the technology of migration itself — it is the operational and process of migration around it."

CISO Forum: What does a realistic IT/OT security architecture look like today — and how does Precision AI change threat detection in OT environments where patching is impossible?

SWAPNA BAPAT: A realistic IT/OT security architecture starts with accepting that OT cannot be treated like enterprise IT. In manufacturing and critical infrastructure, uptime and safety come first, so the priority is visibility, segmentation, controlled access, and continuous monitoring — not simply pushing patches or replacing legacy systems.

The first step is to identify the assets across plants, production lines, remote sites, and connected industrial systems. From there, organizations need to segment critical environments, enforce least-privilege access, monitor traffic between IT and OT, and build response plans that do not disrupt operations.

In OT environments, many systems are too old, too sensitive, or too critical to take offline regularly. AI-led detection can help identify abnormal device behavior, risky communication patterns, exposed assets, and potential exploit attempts earlier.

The goal is not to make OT look like IT. It is to protect industrial environments while respecting operational realities — especially where patching windows is limited or impossible.

CISO Forum: What does a mature supply chain risk program backed by PANW look like — and what is the one supply chain risk Indian CISOs are systematically underestimating in 2026?

SWAPNA BAPAT: A mature supply chain security program today extends well beyond vendor assessments and periodic compliance reviews. Organizations need continuous visibility into SaaS integrations, APIs, managed service providers, software dependencies, cloud environments, and the broader ecosystem of third parties connected to enterprise infrastructure.

One of the most underestimated risks remains at the level of implicit trust that organizations continue to place in interconnected external environments. Many enterprises invest heavily in securing their own infrastructure while assuming partners, vendors, and external platforms to maintain equivalent governance standards and identity controls. Attackers are increasingly exploiting those weakly connected ecosystems to gain indirect access to larger organizations.

The growing dependence on browser-based workflows, machine-to-machine communication, and interconnected SaaS ecosystems is also creating new lateral movement opportunities that many enterprises still struggle to monitor effectively across distributed environments.

Supply chain resilience is rapidly becoming one of the defining cybersecurity leadership priorities for global enterprises. ■

The CISO who codes security into the product before the first line ships

Vineet Daniel, CTO & CISO at PayMe (Huey Tech) on building security into fintech from day one and why AI's pace is the real threat.

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

IN MOST organizations, the CTO and CISO sit on opposite sides of a familiar argument: one wants to move fast, the other wants to move carefully. Vineet Daniel has resolved that tension by becoming both. As CTO and CISO at PayMe, a fintech lending platform navigating the convergence of open APIs, embedded finance, and AI-driven credit underwriting, he has built a security philosophy that treats compliance not as a brake but as a foundation. In this conversation with CISO Forum, Daniel speaks candidly about structuring DPDP consent frameworks without sacrificing product velocity, the three Zero Trust gaps Indian fintechs rarely admit publicly, and why the pace of AI evolution, not any single attack vector, is what keeps him up at night.

CISO Forum: Running both CTO and CISO roles simultaneously — where does the tension between shipping fast and securing deep show up most acutely in a fintech like PayMe (Huey Tech)?
VINEET DANIEL: Being a CTO and a CISO means having two different hats at the same time. One responsibility is to deliver on time using innovative methods and technologies; the other is to ensure everything is compliant and secure. But at the same time, being a CISO doesn't mean the CTO's role has to stop innovating or shipping fast. Both work in conjunction and, in fact, have helped me build security from day one. So whatever we are doing, security is baked in, compliance is baked in, and we always make sure the team is trained and aware of the compliance requirements and regulations set by the RBI. Being both CTO and CISO has, in fact, helped both our team and our mission by enabling us to ship fast, innovate, and deliver secure solutions.

> "Zero Trust is easy to declare for new systems and very hard to retrofit onto what we already have, especially legacy systems."
>
> **Vineet Daniel**
> CTO & CISO, PayMe

CISO Forum: AI-powered credit underwriting relies heavily on alternative data. What are the biggest cybersecurity and data integrity risks that CISOs in lending platforms often underestimate?
VINEET DANIEL: I wouldn't say that CISOs often underestimate anything; that would be an understatement. As a CISO, we have to make sure that, especially after DPDP, PII has come into the picture and into force, the personal data of the user, the ultimate consumer of the platform, remains secure and doesn't get hacked, breached, or used by malicious actors. So, in that sense, whatever we are building, as I said earlier, whether it is on the database, the platform, or in transit, we make sure that it is secure. Even in our models, we have tested guardrails in place. Whenever PII is detected, and someone has inadvertently supplied personal information, a notification is sent. We make sure that everything, especially PII, doesn't enter the public domain, LLMs, or unsecured or unsafe zones.

CISO Forum: With enforcement of the DPDP Act tightening in India, how are you structuring your DPIA and consent management frameworks without slowing product velocity?
VINEET DANIEL: With the DPDP Act now in force, privacy has actually moved from being a good practice to a statutory obligation for vendors. We have built upon the concept of purpose limitation, for example. And it's not just us; all the vendors who are supplying or helping us with data are equally aware. We also met with them while the DPDP framework was being finalized and before it came into force to discuss how we would manage user data. Because in certain places we rely on third-party APIs to obtain data, we have ensured, in collaboration with the vendors, that proper user consent is in place and is timestamped and recorded with audit trails and logs. The purpose is clearly mentioned, visible to the user, and visible to us. By that, we mean that the use of that data is solely for the stated purpose. For example, if we are pulling a bureau report for a particular loan, we use it only for decision-making and then discard it. Whenever that user applies for a new loan, we fetch fresh records.

CISO Forum: What does PayMe's AI-powered fraud detection architecture look like, and how do you balance anomaly detection sensitivity against false positives that could hurt genuine borrowers?
VINEET DANIEL: When it comes to technology, it's not perfect. There are certain flaws and false positives. It depends on the ratio at which they occur, but thankfully, you can keep tweaking the percentage or threshold of what should be considered the correct parameter. For example, if a user uploads a live selfie, there is another check we have in place that brings AI into the picture. We verify that it is a live selfie, not a recorded one or

a reused image. There are certain checks that we put in place.

Now, there can be false positives, but we also have to consider the impact of real fraud being ignored just because we don't want users to be flagged by the AI system. So there is a small percentage we consider acceptable, and in those cases, we direct them for manual review rather than discard them.

We send them for manual review, and the feedback is fed back into the system. If the system's flag was correct or incorrect, learning happens through reinforcement learning, and we continuously improve those systems. So yes, there could be false positives, but it is always a learning system. While we try to reduce false positives, people with malicious intent keep devising innovative ways to bypass controls. So it's always an ongoing learning process.

CISO Forum: As embedded finance expands through open APIs, the attack surface grows dramatically. What's your security philosophy for third-party API integrations with institutional partners?

VINEET DANIEL: Embedded finance is real, and it went through a high-growth phase in which unit economics held up.

What I have learned is that embedded lending succeeds on two things: integration quality and absolute clarity of where the credit risk sits. As I mentioned earlier, we had many meetings with vendors, especially in embedded finance, where account aggregators play a key role. They have their own consent management systems and their own security practices in place. The DPDP Act does not apply just to us; it is equally applicable to our vendors. They are equally responsible and aware that they must maintain compliance.

So it's not just us who need to be compliant; they do, too. When we onboard any vendor, we review each other's compliance posture. We have a checklist that is followed with every vendor onboarding. We also have ISO and PCI DSS automation, which strengthens our security posture and gives our customers and vendors confidence that we are not creating security risks.

CISO Forum: CISOs across BFSI are nervous about GenAI. At PayMe, are you deploying it more aggressively as a capability, or are you still focused on managing it as a risk?

VINEET DANIEL: I am enthusiastic about Generative AI and deliberately conservative about where it goes. Where it has delivered real value for us is in engineering productivity, internal knowledge management, operations, and, increasingly, customer support. We have deployed a customer support bot that has offloaded 60–70% of our customer queries from the support team to the bot.

The other area where we focus is governance. We authorize AI access through a controlled environment and use hyperscalers to do so. We don't allow unmanaged personal AI tools on corporate devices or within office premises. Everything is done through enterprise accounts. We have measures in place to handle sensitive data.

I wouldn't call it anti-AI. In fact, it's the opposite. It's about treating GenAI as something that must operate with clear access controls, audit trails, and accountability. That's what lets you adopt it confidently rather than nervously.

So I wouldn't say that, as a CISO, we need to be nervous or overly conservative about its adoption. This is a change that we cannot resist; we have to adapt to it sooner or later. At PayMe, we have already started adopting it, and the biggest change we have seen is in engineering and other departments, where it has increased efficiency multifold. It is another very efficient tool in our toolkit.

CISO Forum: SaaS-based lending platforms mean deep dependency on third-party vendors. How do you approach vendor risk management and security due diligence at scale?

VINEET DANIEL: We have our own checklist that we have created over time. It has gone through many changes and iterations during this period.

For example, with DPDP coming into the picture, we made some changes again and added many clauses to it. We ensure that most of the vendors we onboard have VAPTs, relevant certifications, and proper governance practices within their organizations. We assess how seriously they take DPDP compliance and review the frameworks they use. We also conduct a small POC before onboarding and evaluate all the good practices they follow.

We ensure that the checklist is at least 90–95 percent green. At times, certain items may still be in progress. For example, one of our vendors was in the process of obtaining PCI DSS certification for their operations, so we gave them an exception with a time-bound limitation. We agreed that by a specific date, they would have to share the certification with us. So, this is how we go about it.

There is no casualness in due diligence. We ensure that everything, especially PII and customer data, is secured and transferred securely.

There are measures and validations at each and every point where the data is transferred or changes hands.

CISO Forum: Has PayMe moved to Zero Trust architecture? What are the practical implementation gaps that Indian fintech CISOs rarely discuss publicly?

VINEET DANIEL: There are certain things that Indian fintech CISOs do not want to discuss publicly, and of course, for obvious reasons. But yes, we've had close discussions among ourselves and with peers from different fintechs about how they are handling security, and ZTNA is one of the very good tools that helps us achieve that. This is also one of the questions I have been asked frequently in the past.

The first thing nobody likes to admit is that Zero Trust is actually easy to declare for new systems and very hard to retrofit onto what we already have, especially legacy systems. Most Indian fintechs, including ours, have legacy services and integrations built before Zero Trust was implemented. We are actively working on it and building new architectures in which Zero Trust principles are implemented from the start. We are also in talks with several vendors at the moment, and by next month, we will fully implement Zero Trust in our architecture.

The second gap is identity. Zero Trust actually lives or dies on strong identity management for every user, service, laptop, desktop, or any other machine connected to the network. Human users may have MFA, but service accounts and emergency access mechanisms can sometimes accumulate privileges over time because tightening them may risk breaking something in production.

"The pace of evolution of AI systems, the skills required to keep up, and the resources needed to defend against increasingly sophisticated threats — those are the things that create sleepless nights."

So, we have to be careful about that as well. The new changes we are bringing in cannot apply only to newer systems, while legacy systems remain unchanged. That gap needs to be plugged before production deployment, which increases implementation time.

The third gap that I can think of at the moment is third parties, especially in a regulated lending ecosystem. We operate in an ecosystem of partner integrations and embedded finance, which we discussed earlier. A Zero Trust architecture or posture effectively stops at your own boundary.

The moment data crosses over to a partner, the situation changes. Data may be fully secured while it is with us, with all the required controls and safeguards in place. But once it passes on to a partner, we have to trust their network and trust their security maturity. We can define requirements contractually and include them in legal agreements, but it is difficult to verify and validate them regularly.

These are the three gaps that I can think of right now, and the third one is probably the most important.

CISO Forum: Cybersecurity talent in India is scarce, especially at the intersection of fintech and AI security. How are you building and retaining that capability internally at PayMe?

VINEET DANIEL: As a CISO at an NBFC like PayMe, the threat model extends beyond financial inclusion. Modern identity abuse is also part of it. KYC fraud, account takeover, and API abuse are all concerns. For example, we have had instances where a person used his father's phone and applied for a loan without the father even being aware of it.

So, I'll give you a very honest answer. AI lowers the cost of attacks—deepfakes, KYC fraud, phishing at scale, and automated browsing attacks—but it is also a practical way to defend at that volume. No human team can review fraud signals quickly enough.

We need AI systems that work at scale and can operate in real time. At the same time, the importance of human cybersecurity personnel does not decrease; it actually increases. AI, if we think of it as an enabler, is also a threat because malicious actors may use it for purposes other than its intended use.

That is where cybersecurity personnel come into the picture. They continuously track these changes, test AI models, and identify weaknesses and risks within them.

Coming back to your question about the shortage of cybersecurity personnel, the shortage is not as severe if you can find a trusted partner to whom you can outsource parts of your cybersecurity workload. For example, you can have a cybersecurity professional within the company who oversees cybersecurity partners and vendors, and that can help resolve many of the challenges. ■

Half your AI security controls are failing

AI-powered attacks, invisible machine identities, geopolitical risks, quantum threats, and vulnerable supply chains are redefining cybersecurity, making it a core boardroom priority and a critical business survival issue.

By **CISO Forum** | editor.tech@timesgroup.com

ARTIFICIAL INTELLIGENCE is no longer a tech experiment; it's the backbone of how modern businesses operate. According to Proofpoint's 2026 AI and Human Risk Landscape report, a staggering 87% of organizations have AI assistants deployed well beyond the pilot stage, and 76% are actively rolling out autonomous agents, AI systems that can independently plan and execute tasks without human approval at each step.

The problem? Security didn't get the memo. Only 48% of organizations say security was embedded in their AI strategy from the beginning. The remaining 52% admit security is either playing catch-up, inconsistent, or entirely reactive. AI was deployed to production before the guardrails were put in place.

Controls exist. They don't work.

Here's where things get alarming. A confident 63% of organizations say they have AI security controls in place. That sounds reassuring — until you dig into what those controls are actually doing.

More than half (52%) of organizations are not fully confident that their controls would detect a compromised AI. And the numbers back up that fear: among organizations that have security controls deployed, 50% still reported a suspicious or confirmed AI-related incident. That's not a minor gap in coverage. That's a coin flip.

The report's conclusion is blunt — coverage is being mistaken for control.

Threats are everywhere, not just in email

When organizations that experienced AI-related incidents were asked where threats appeared, the results showed no safe harbor. Email led to the list at 67%, followed by SaaS and cloud apps at 57%, AI assistants and agents at 53%, and collaboration tools, file-sharing, and social platforms at 49%.

This matters because AI is deeply embedded in these channels. Companies are using AI for customer support (69%), internal chat summarization in Slack or Teams (67%), and email drafting (63%). Attackers, the report notes, follow the operating model. Where AI works, threats follow.

"More than half (52%) of organizations are not fully confident that their controls would detect a compromised AI."

When something goes wrong, nobody can reconstruct what happened

The most chilling finding in the report is about incident response. Only one in three organizations say they are fully prepared to investigate an AI- or agent-related incident. Nearly 95% report that managing multiple, disconnected security tools is at least moderately challenging — and 53% call it very challenging.

The consequence is structural: 41% of organizations cannot correlate threats across multiple channels. When an attacker enters via email, escalates through a collaboration tool, and exfiltrates data via an AI integration, siloed tools cannot reconstruct the chain of events.

The industry is waking up - slowly

Organizations aren't standing still. 61% plan to expand AI protections, 56% aim to extend collaboration coverage, and 53% intend to move to a unified security platform—only 3.9% plan to keep the status quo.

The Proofpoint report's message is clear: the organizations that will successfully scale AI aren't those deploying the most tools — they're those that build security around how modern work happens, across people, platforms, suppliers, and AI systems alike. Anything less is a gamble that most organizations are already losing.

The 2026 AI and Human Risk Landscape report is based on a survey of 1,453 security professionals conducted in January 2026, spanning 20 industries across 12 countries. ■

AI's next chapter: What a UK security report reveals

AI's next wave won't come from one breakthrough; diversity, agility, and governance will define what's ahead.

By **CISO Forum** | editor.tech@timesgroup.com

A NEW report from the Alan Turing Institute's Center for Emerging Technology and Security (CETaS), The Next Frontier: Security Implications of Future AI Paradigms, offers one of the most comprehensive assessments yet of where AI is heading and what that means for governments, businesses, and national security.

The current AI boom isn't over, but It's changing

The transformer-based AI models powering today's chatbots and coding tools still have room to grow but not in the way we've grown accustomed to. Simply throwing more computing power at existing systems is hitting a wall. High-quality training data is running out, and the energy demands are staggering next-generation AI data centers could require gigawatt-scale power, roughly the output of several nuclear reactors.

Instead, the report argues, future gains will come from a combination of smarter techniques of reinforcement learning, better reasoning at inference time, and more sophisticated system design rather than brute-force scaling alone.

A more fragmented AI landscape is coming

Rather than a single dominant breakthrough, the report maps out 15 distinct AI paradigms likely to shape the next decade. These range from agentic AI systems models that can autonomously plan, use tools, and take actions to neuromorphic computing, world models, and quantum machine learning.

Some, such as agentic systems and lean, efficient models, are already commercially deployed. Others, such as thermodynamic computing and quantum machine learning, remain in their early stages. The key takeaway: the future AI landscape will be more diverse, harder to govern, and harder to predict.

Agentic AI: Powerful but dangerous

The report flags agentic AI as both the most commercially exciting and most security-concerning near-term development. When AI agents gain access to private data, interact with untrusted content, and can take real-world actions such as transferring money or accessing systems, the attack surface expands dramatically. Errors compound, accountability blurs, and oversight frameworks struggle to keep pace.

"The future AI landscape may become more varied in architecture, training methods, deployment models, and hardware base."

The geopolitical stakes

The US and China together employ 70% of the world's top machine learning researchers and control 90% of AI training compute. China is aggressively pursuing open-source AI dominance, with its models now downloaded more frequently than American counterparts. The report warns that asymmetric AI capabilities between rival powers could destabilize deterrence and lower the threshold for conflict.

What the UK and others must do

For nations outside the US-China duopoly, the report's prescription is strategic agility over frontier ambition: invest in a deep skills base, build fine-tuning infrastructure, and focus on deploying AI effectively at scale rather than racing to build the largest model.

The message for policymakers worldwide is clear — prepare for many possible AI futures, not just one.

Cybercriminals are already playing the World Cup

FortiGuard Labs' FIFA World Cup 2026 Cyberthreat Landscape Report reveals a sophisticated, coordinated criminal operation already in motion.

By **CISO Forum** | editor.tech@timesgroup.com

BEFORE THE first whistle blows at FIFA World Cup 2026, cybercriminals have already kicked off their own tournament, and the stakes are real money, stolen identities, and compromised devices.

Fortinet's FortiGuard Labs has published a detailed threat intelligence report mapping the scale of digital fraud building around the world's most-watched sporting event. The findings are striking: over 13,000 FIFA-themed domains registered in just five months, more than 1,700 fake social media accounts, and hundreds of thousands of credentials already circulating in underground markets.

A domain surge that signals intent

Between January and May 2026, cybercriminals registered thousands of FIFA-lookalike domains at an accelerating pace — from 235 in January to nearly 5,000 in April alone. Roughly 8.8% were classified as outright malicious. Most mimicked ticketing portals, streaming services, or hospitality platforms. The surge isn't coincidental; it reflects deliberate infrastructure buildout ahead of the tournament.

Fake tickets, real losses

Ticket fraud is the most visible vector for scams. Researchers identified fully operational fake ticketing sites, including one registered just weeks before the tournament, that replicated FIFA's official branding down to the checkout page, capturing billing details, card numbers, and login credentials. Underground forums and Telegram channels are advertising discounted tickets bundled with fraudulent flight and hotel packages, often demanding payment via cryptocurrency or wire transfer to avoid traceability.

The job scam no one saw coming

One of the report's more alarming findings involves a coordinated job-recruitment phishing campaign. Fraudulent websites impersonating FIFA and its corporate sponsors, including major hospitality and beverage brands, lured victims with fake calendar meeting invites, then directed them to credential-harvesting pages that mimicked Google login. A

"Before the first whistle blows at FIFA World Cup 2026, cybercriminals have already kicked off their own tournament, and the stakes are real money, stolen identities, and compromised devices."

shared Google Analytics ID across dozens of these domains points to a single organized threat actor running the operation at scale. a project IRR of 20.3% and an equity IRR of 28.4%.

Malware in the streaming queue

Fake streaming platforms and trojanized betting apps are another growing risk. A malicious executable disguised as a popular betting application was found to deploy ransomware-linked encryption techniques and communicate with external servers via legitimate cloud platforms to evade detection.

Credentials already on the dark web

The report found over 270,000 credentials from fans visiting FIFA-related websites already present in stealer log datasets, harvested by malware families including Vidar, LummaC2, and RedLine. More than 260 credentials tied to FIFA employees were also identified.

What this means

The report makes clear this isn't opportunistic cybercrime — it's organized, infrastructure-backed fraud. For fans, the message is straightforward: buy only through official channels, verify every URL, and treat unsolicited job offers or ticket deals with serious skepticism. ■

Your inbox is a battlefield, and you're losing

Barracuda's 2026 report reveals one in three emails is malicious or spam, as AI and Phishing-as-a-Service make attacks cheaper and harder to detect.

By **CISO Forum** | editor.tech@timesgroup.com

EVERY DAY, billions of emails cross the internet. And according to Barracuda's 2026 Email Threats Report, a staggering one in three of those messages is either malicious or unwanted spam — up sharply from one in four just last year. In January 2026, Barracuda Research analyzed over 3.1 billion emails, and what they found should alarm every executive, IT manager, and employee who has ever clicked a link at work.

Phishing still reigns supreme

Nearly half — 48% — of all malicious email activity is phishing. These aren't clumsy, poorly written scams anymore. Today's phishing emails are polished, personalized, and often indistinguishable from legitimate business communications. Attackers impersonate trusted brands like Microsoft, DocuSign, and SharePoint, luring employees into surrendering login credentials that can unlock entire corporate networks.

The rise of "crime-as-a-subscription"

Perhaps the report's most alarming revelation is the explosion of Phishing-as-a-Service (PhaaS). In this criminal business model, would-be hackers subscribe to ready-made attack toolkits, complete with fake login pages, automation, and hosting. A dramatic 90% of high-volume phishing campaigns in 2025 used PhaaS kits a massive jump from just 30% in 2024. The barrier to entry for launching a sophisticated cyberattack has essentially collapsed. Even technically unskilled criminals can now run large-scale campaigns targeting thousands of victims simultaneously.

Many of these kits also come with MFA bypass tools, meaning that two-factor authentication long considered a gold-standard defense is no longer enough on its own.

Your files are being weaponized

Attachments remain a favored attack vehicle, but the methods have grown more cunning. More than 10% of all HTML attachments are malicious — the most weaponized file type by a wide margin. When opened, these files render in a browser and silently redirect users to credential-harvesting websites.

Even more concerning: 70% of malicious PDFs now contain QR codes leading to phishing websites. By embedding a QR code inside a trusted-looking document, attackers shift the attack to a victim's mobile phone — a device that typically sits outside a company's security perimeter and monitoring systems.

"Email is no longer just a communication tool. It is the primary gateway through which cybercriminals enter, and the sophistication of those attempting entry has never been higher."

One-third of companies get hacked every month

Account takeover where attackers access a real employee's email and operate from within is no longer rare. 34% of companies experience at least one incident each month. Attackers don't just steal data; they manipulate inbox rules, forward sensitive emails externally, and launch phishing from within, using the victim's trusted identity. A quarter of incidents involved suspicious changes to inbox rules, helping attackers stay hidden longer.

Links are the new weapon of choice

As companies improve attachment scanning, attackers adapt. Barracuda Research has identified a marked increase in URL-based attacks over direct file attachments. Criminals now host malicious content on reputable platforms like SharePoint and Google Drive, making links appear legitimate even programming them to appear harmless during security scans, only activating once an email reaches a real inbox.

What businesses must do now

The report's prescription is clear: layered defense is no longer optional. Organizations must invest in AI-enhanced email filtering, zero-trust access controls, automated incident response, and regular employee training, including simulated phishing tests. Backing up data and auditing configurations are equally non-negotiable. Email is now the primary gateway through which cybercriminals enter, and their sophistication has never been higher. ■

India's data centers: Powering the AI race

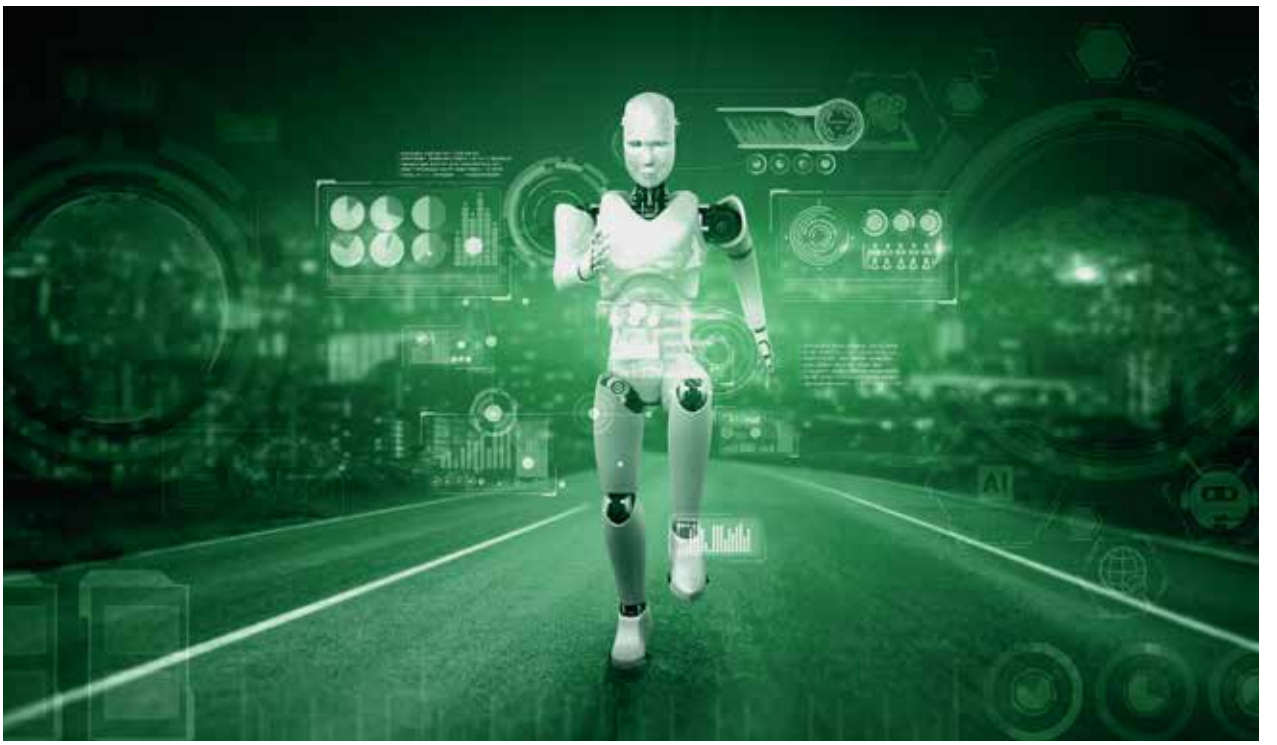

India's data center capacity is set to triple by 2030, fueled by surging AI adoption and $25 billion in investment.

By **CISO Forum** | editor.tech@timesgroup.com

INDIA IS quietly building one of the world's most ambitious digital infrastructure stories. As AI moves from boardroom buzzword to operational reality, the country's data center industry is scaling at a pace few anticipated and the numbers demand attention.

A market on the move

India's installed data center capacity stood at 1.6GW in December 2025 and is projected to hit 5GW by 2030, a 26% CAGR. To get there, developers are sitting on an active pipeline of over 3 GW, requiring approximately USD 25 billion in capital investment. This isn't speculative; it's already under construction.

The AI engine behind it all

India's domestic AI market, currently valued at USD 13 billion, is forecast to reach USD 131 billion by 2032, growing at 39% annually. Over 80% of Indian enterprises are prioritizing AI adoption, and 45% are already implementing it. The IndiaAI mission launched in March 2024 with a USD 125 billion budget has secured commitments for 38,000+ GPUs, which alone will require 48MW of AI-ready data center capacity.

GPUs: The new gold

The Avendus report makes clear that GPUs are the engine of this AI infrastructure race. A single ChatGPT query requires over 30,000 Nvidia A100 GPUs running simultaneously. Training a foundational model takes 12–15 months of continuous GPU compute. NVIDIA commands 90–95% of the GPU market. For investors, a 3,000-GPU cluster deployment is projected to deliver a project IRR of 20.3% and an equity IRR of 28.4%.

Where the money is going

Mumbai dominates, accounting for roughly 50% of India's current data center capacity. Chennai ranks second, benefiting from five submarine cable landing stations. Hyderabad is emerging fast, attracting both colocation players and hyperscaler self-builds.

"A single ChatGPT query requires over 30,000 Nvidia A100 GPUs running simultaneously."

The industry has seen over USD 5 billion in transaction activity in just three years, with deals such as TCS-TPG's USD 2 billion HyperVault commitment and Nxtra's USD 1 billion fundraise signaling strong institutional confidence. .

Valuations reflect the urgency

Globally, data center assets are trading at 20–30x EV/EBITDA multiples. In India, listed players like E2E Networks command multiples above 44x. Several domestic operators — Sify, Yotta, ESDS — are lining up IPOs, signaling that public markets are ready to price this sector at a premium.

The bottom line

India's data center buildout is no longer a supporting act for the digital economy. It is the digital economy's backbone — and increasingly, the infrastructure on which India's AI ambitions will either stand or fall.

Based on the Avendus Capital report "Data Centers: Powering India's AI Boom" (May 2026).■

India's banks are losing the AI cyber arms race

AI has slashed exploit times by 94%, but Indian banks' cyber defenses and confidence haven't kept pace.

By **CISO Forum** | editor.tech@timesgroup.com

A NEW report from Boston Consulting Group and the Data Security Council of India lays bare an uncomfortable truth: artificial intelligence has given attackers a speed advantage that Indian banks, insurers, and financial firms cannot yet match.

Attacks are getting cheaper, faster, and harder to stop

The numbers tell the story. The time it takes hackers to exploit a newly disclosed vulnerability has collapsed from 745 days in 2020 to just 44 days today — a 94% drop. The cost of mounting a sophisticated attack has fallen by more than 70%, with frontier AI tools now able to attempt a full network intrusion for as little as $80. India's financial sector is being hit 1.6 times more intensely than the global average, and cyber incidents handled by India's national response agency have more than doubled since 2021, from 1.4 million to 2.9 million.

The confidence gap

Despite the alarm, preparedness hasn't kept pace. In a survey of over 40 Indian chief information security officers, not a single control area — from identity management to forensic response — crossed the 50% confidence threshold for withstanding an AI-powered breach. Meanwhile, 43% of CISOs admit attackers are already outrunning their defenses, yet only 19% have raised cyber budgets by more than 10% to respond. Most incremental spending is going toward AI tools, but overall investment as a share of IT budgets still trails global peers.

A call for "Synchronized" defense

Rather than prescribing more tools, the report argues Indian BFSI needs a fundamental shift in how cybersecurity is organized. It outlines five fronts institutions must align: tying cyber investment to actual

"Just 29% of Indian financial institutions have both a designated AI security owner and a formal policy in place."

business risk, breaking down silos between IT, risk, legal and business teams, treating vendors as extensions of the enterprise rather than annual compliance checkboxes, unifying defenses against insider risk and customer-facing fraud, and — critically — sharing threat intelligence across the industry rather than fighting attackers alone.

The bottom line

The report's authors, including BCG's Nisha Bachani and DSCI's Vinayak Godse, frame the next 12 to 18 months as decisive. Institutions that move now to build governance and cross-functional coordination, rather than simply buying more security software, will be the ones that convert resilience into a genuine competitive edge. Those that don't, the report warns, will remain permanently on the back foot — defending yesterday's perimeter against tomorrow's threat.
Source: "Cybersecurity in the Age of AI," *Boston Consulting Group and Data Security Council of India, May 2026.* ■

The basic problem remains: human error

HALF OF 2026 is already behind us, and despite the relentless pace of technological change, one reality remains unchanged: people continue to be the weakest link in cybersecurity.

The challenges facing today's CISOs have only multiplied. Geopolitical tensions are reshaping cyber risk, an evolving regulatory landscape is demanding greater compliance, and AI is transforming both attack and defense at unprecedented speed. Every week brings a new AI-driven threat, a sophisticated phishing campaign, or another debate about autonomous security. Yet under all the complexity lies a surprisingly familiar problem, human error.

Conversations with CISOs across India consistently point to the same conclusion. The majority of security incidents still start with a simple mistake: an employee clicking a malicious link, sharing sensitive information with the wrong recipient, using weak credentials, or bypassing established security processes in the name of convenience. Even organizations that invest heavily in advanced security technologies remain vulnerable when everyday security habits fail.

This reality has pushed governance to the forefront. Security can no longer depend on deploying better tools; it must be embedded into business processes, decision-making, and organizational culture.

The rapid adoption of generative AI has only reinforced this need. Gartner predicts that Shadow AI is inevitable. Rather than trying to eliminate it through restrictive, centralized controls, organizations should adopt collaborative governance models that encourage business accountability. By monitoring usage patterns, understanding exceptions, and co-creating practical guardrails with employees, security leaders can regain visibility while reducing the risks associated with sensitive data being shared with AI systems.

Equally important is strengthening security behavior and culture. Gartner advises cybersecurity leaders to look beyond technical controls and invest in security behavior and culture programs (SBCPs) that encourage safer GenAI practices across the organization.
For all the advances in AI, zero trust, and automation, cybersecurity still comes down to one timeless truth: technology can reduce risk, but only people can truly manage it. That is where the next generation of cyber resilience will be won or lost. ■

"Technology can reduce risk, but only people can truly manage it."

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

AMD

presents

ET Edge | 27th Annual Conference

CIO&LEADER

The Agentic Enterprise

PLATFORMS. PEOPLE. POLICY. PROFITS.

co-presented by

NxtGen

powered by workday.

co-powered by outsystems

#CIOandLeaderConference

30 JULY – 02 AUGUST 2026

FAIRMONT, JAIPUR

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

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

WHAT's IN FOR YOU

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

EMINENT SPEAKERS

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

Rana Ashutosh Kumar Singh
MD (International Banking, Global Markets & Technology), SBI

Madan Sunder Das
Monk, ISKON, Founder, Spiritual Guide and Community Leader, EVOLVE Pune

Be Where India's Tech Leaders Meet.

PRESENTING PARTNER: AMD | CO-PRESENTING PARTNER: NxtGen | POWERED BY PARTNER: workday. | CO-POWERED BY PARTNER: outsystems | GOLD PARTNER: dynatrace | Commvault | ASSOCIATE PARTNERS: CROWDSTRIKE | DELTA | TECHNO DIGITAL

EXHIBIT PARTNERS: BARCO ClickShare | Barracuda | InfoBeans | ivanti | Newland AIDC | PrivacyPillar

SKILLING PARTNER: TIMESPRO | EDUCATION PARTNER: BENNETT UNIVERSITY | BUSINESS INSIGHTS PARTNER: INSIGHTS | CONCEPTUALIZED AND EXECUTED BY: ET | CIO&LEADER | NEXT100

For Sponsorship, Write to

Naveen Singh
c-naveen.singh@timesgroup.com
9901300772

Hafeez Shaikh
hafeez.shaikh@timesgroup.com
9833103611

Aanchal Gupta
aanchal.gupta1@timesgroup.com
9651841119

Supriya Sahoo
supriya.sahoo@timesgroup.com
8095056886

NxtGen

presents

ET Edge

CISO FORUM MID YEAR CONFERENCE '26

Building A Cyber Resilient Enterprise

12 -14 June 2026
Hayatt Regency Jaipur Mansarovar

Thank You

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

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

PRESENTING PARTNER: NxtGen

GOLD PARTNER: ARMIS from ServiceNow

SILVER PARTNERS: airtel, netskope, proofpoint, CYBERPWN, rubrik

STRATEGIC CASE STUDY PARTNER: SECLORE

ASSOCIATE PARTNER: orca SECURITY

EXHIBIT PARTNERS: Barracuda, CyberAssure, cybermindr, DIEVAS, secpod, HarkX, MATTERS, CyberAssure, NeoSOFT, SyskeyOT, techza, Silverfort

SKILLING PARTNER: TIMESPRO

EDUCATION PARTNER: BENNETT UNIVERSITY, THE TIMES GROUP

BUSINESS INSIGHTS PARTNER: ET INSIGHTS

CONCEPTUALIZED AND EXECUTED BY: ET Edge | CISOFORUM