JULY-AUGUST 2026
SUPPLEMENT OF CIO&LEADER

A 9.9 GROUP PUBLICATION
www.cisoforum.in

CISOFORUM

Security For Growth And Governance

Who Owns Your AI Data?

DPDP is forcing CISOs to confront privacy risks spreading across the enterprise PG.09

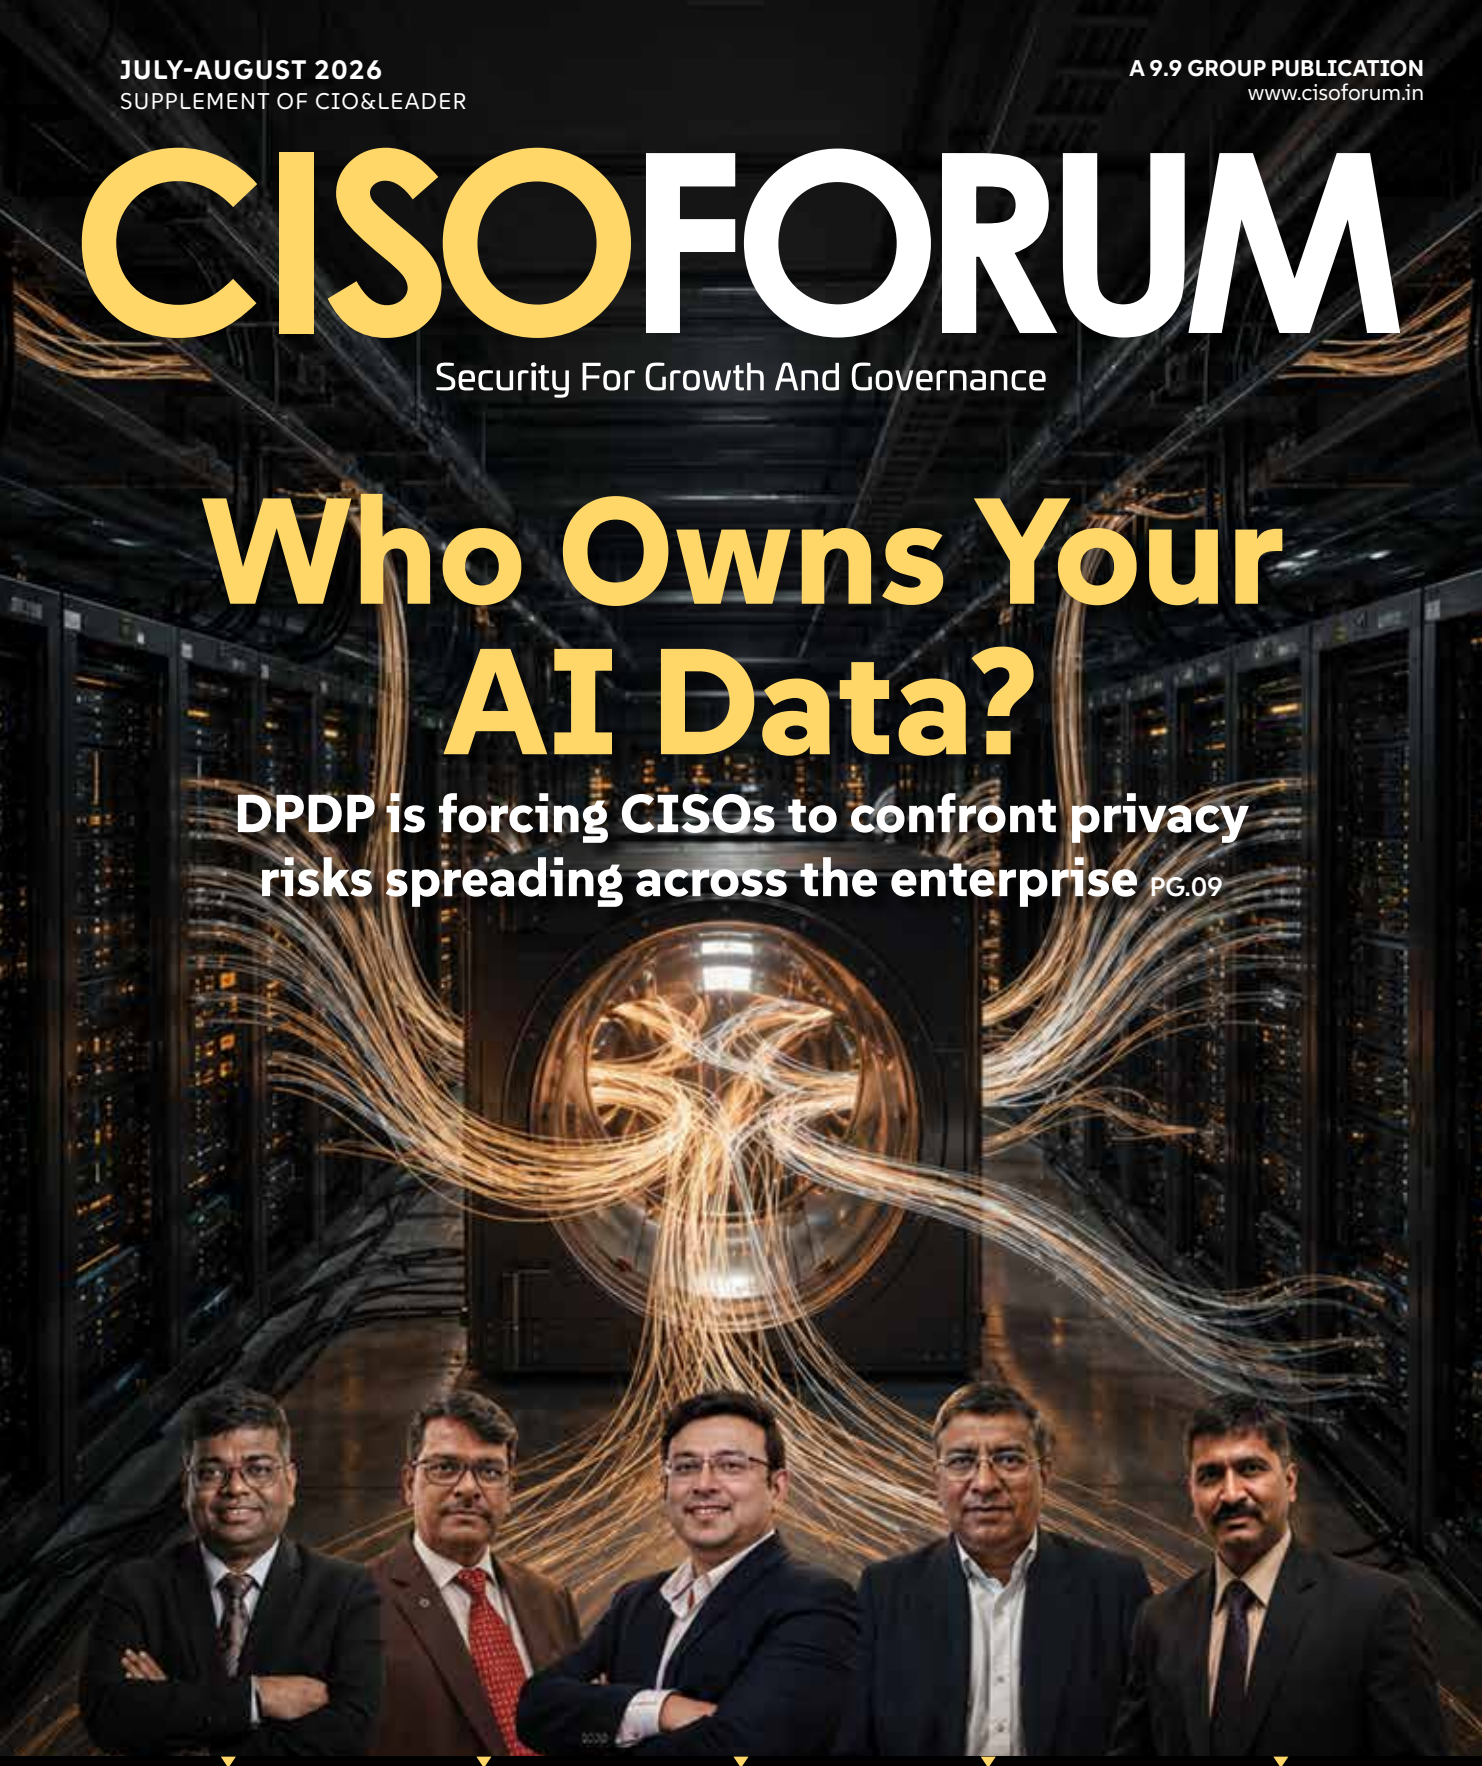

Prakash Kumar Ranjan
CISO & DPO,
Sammaan Capital

Shermugadurai C
VP & CISO,
Tamilnad Mercantile Bank

Tarovar Verma
Head of Enterprise Security,
PhysicsWallah

Vishwas Pitre
CISO & DPO,
Zensar Technologies

Harsha Vardhan Jonnalagadda
CISO & DPO,
Birlasoft

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

freshworks outsystems

#CIOandLeaderConference

30 JULY – 02 AUGUST 2026

FAIRMONT, JAIPUR

THE AGENTIC ENTERPRISE

PLATFORMS. PEOPLE. POLICY. PROFITS.

THANK YOU FOR MAKING IT EXTRAORDINARY.

The **27th Annual CIO&Leader Conference** brought together India's top CIOs & IT leaders for four unforgettable days of bold conversations, transformative ideas and meaningful connections.

218+ CIOs & IT Leaders | 4 Days of Experience

From **Platforms. People. Policy. Profits.** to the possibilities of **the Agentic Enterprise,** the conversations have only just begun.

To every CIO, IT Leader, speaker and partner —
Thank you for being part of the journey.

PRESENTING PARTNER: AMD

CO-PRESENTING PARTNER: NxtGen

POWERED BY PARTNER: workday.

CO-POWERED BY PARTNERS: freshworks, outsystems

GOLD PARTNERS: dynatrace, HPE, TATA COMMUNICATIONS, snowflake

SILVER PARTNER: STTelemedia Global Data Centres

EXPERIENCE PARTNER: synerome rewedge

ASSOCIATE PARTNERS: AUTODESK, Commvault, CROWDSTRIKE, DELTA, digitaltrack, GENESYS, InfoBeans, servicenow, Liferay, NOPALCYBER, ORACLE, Procol.AI, Prophaze, BD SOFTWARE DISTRIBUTION PVT LTD, rubrik, savex, SOLARWINDS, TECHNO DIGITAL, YOTTA

EXHIBIT PARTNERS: BARCO ClickShare, Barracuda, CyberAssure, VARONIS, Forcepoint, ivanti, NeoSOFT, Newland AIDC, nividous, PrivacyPillar, RSA, Cotelligent

SUPPORTING PARTNER: OneDios

EVENT TECH PARTNER: CONNECT, ZEBIQ

SKILLING PARTNER: TIMESPRO

EDUCATION PARTNER: BENNETT UNIVERSITY THE TIMES GROUP

BUSINESS INSIGHTS PARTNER: ET INSIGHTS

CONCEPTUALIZED AND EXECUTED BY: ET Edge, CIO&LEADER

Privacy on the Clock

Virtually every large Indian enterprise is now aware that data protection has legal weight, needs board attention, and has financial consequence. But what separates the prepared from the exposed is execution.

Breach notification is the most unforgiving of DPDP's obligations. The Act gives fiduciaries 72 hours to report a breach to the Data Protection Board after 13 May 2027. There is no grace period. However, many Indian enterprises are still running this process manually, without the detection tooling, escalation runbooks, or rehearsed tabletop exercises that a hard deadline demands. While a policy document can be drafted in a week, a tested breach-response capability requires time that gets scarcer as 2027 approaches.

Consent management carries a nearer deadline and a harder technical problem. From 13 November 2026, Consent Managers will begin registering with the regulator, and enterprise systems will need to interoperate with them. This is not about issuing a compliant notice, but actually exchange consent state with a third-party platform. Most organizations are still at the notice-redesign stage, which solves the legal requirement but not the integration one.

Underneath both sits a classification problem that predates any technology build: knowing whether your organization is a Significant Data Fiduciary. Enterprises unsure of their own status are, almost by definition, unable to size a DPIA program, a DPO mandate, or an audit cadence. This is arguably the single cheapest fix available, a paperwork exercise rather than an infrastructure one. But work needs to start immediately.

Governance of ownership matters more than most CISOs assume. A dedicated privacy owner — DPO or CPO — consistently outperforms even well-intentioned joint ownership models spanning CISO, Legal, and Business, where accountability tends to diffuse rather than sharpen. And manufacturing OT-heavy data estates deserve a separate budget line altogether; factory-floor systems were never designed for the access controls that DPDP now assumes.

Implementing privacy is systems engineering — not a policy exercise. The enterprises that internalize that distinction now will spend 2027 on gaining an advantage, not performing remediation. ■

"An effective breach-response capability cannot be assembled after an incident."

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

CONTENTS JULY-AUGUST 2026

COVER STORY

Cover Design by:
Shokeen Saifi

Please Recycle This Magazine And Remove Inserts Before Recycling

NEWS & VIEWS

INTERVIEW

SPECIAL FEATURE

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: **Jagrati Rakheja**
Senior Correspondent: **Punam Singh**

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**
Chief Manager- Projects: **Sidharth Ranga**

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**

For editorial queries write to:
editor.tech@timesgroup.com

For sales/business queries write to:
cio.leader@timesgroup.com

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

Editor: **Vikas Gupta**

9.9 GROUP

Sammeer Kulkarni joined Bajaj Life Insurance as CISO

Bajaj Life Insurance has appointed Sammeer Kulkarni as Chief Information Security Officer (CISO). He will lead cybersecurity strategy, information security governance, cyber risk management, compliance, and resilience initiatives. With over 27 years of experience, Kulkarni has held senior security roles at Zurich Kotak General Insurance, BNP Paribas, IndusInd Bank and IndiaFirst Life Insurance.

Sunil Varkey joins Hexaware Technologies as EVP & CISO

Hexaware Technologies has appointed Sunil Varkey as EVP & Chief Information Security Officer (CISO). He will lead to global cybersecurity, resilience, risk management, governance, and compliance. Varkey brings over three decades of experience and joins from TAHAKOM. He previously held senior cybersecurity leadership roles at Forescout, HSBC, Symantec, Wipro and Idea Cellular.

Ashwini Kumar Choudhary joined RBL Bank as Chief Risk Officer

RBL Bank has appointed Ashwini Kumar Choudhary as Chief Risk Officer-Designate. He will strengthen enterprise risk management, credit, market and operational risk oversight, and regulatory compliance. Choudhary joins from Union Bank of India, where he served as Group Chief Risk Officer. He brings nearly three decades of banking and risk management experience.

Subhash Kelkar joined NSDL as Executive Director

NSDL has appointed Subhash Kelkar as Executive Director – Critical Operations & Board Member. He will oversee technology, cybersecurity, resilience, infrastructure, business continuity, and operational risk. Kelkar joins BSE, where he served as CTO. With over three decades of experience across financial services and technology, he has held leadership roles at ICICI Securities, CMS Info Systems, IIFL, and Tech Mahindra.

Vijay Thavasi joined Blackstone as SVP – Cybersecurity Head, APAC

Blackstone has appointed Vijay Thavasi as Senior Vice President – Cybersecurity Head, APAC. He will lead regional cybersecurity strategy, cyber risk, cloud and application security, AI governance, threat intelligence and resilience. Thavasi joins Rakuten, where he served as CISO & IT Head – APAC, bringing over 17 years of cybersecurity experience.

Shrikant Iyer joined Aditya Birla Capital as Head of Cyber Security

Aditya Birla Capital has appointed Shrikant Iyer as Head of Cyber Security. He will lead cybersecurity strategy, governance, cyber risk, compliance, security operations, and resilience across the financial services ecosystem. Iyer previously served as CISO at Aditya Birla Health Insurance and Aditya Birla Insurance Brokers, bringing over 28 years of cybersecurity experience.

CISOs must shift focus as AI reshapes cybersecurity

Gartner urges CISOs to modernize identity, prioritize resilience, and embrace AI-driven innovation to counter evolving threats.

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

Chief information security officers (CISOs) need to rethink identity management, redefine success, and lower barriers to innovation as AI transforms both threats and opportunities, according to Gartner. The insights came from the Gartner Security & Risk Management Summit held June 1-3 in National Harbor, Maryland.

Identity needs a rethink

The explosion of AI agents and machine-to-machine interactions is straining traditional identity and access management systems. Old models built for human users with fixed roles can't keep up with millions of machine identities operating at once. Gartner predicts that by 2028, 25% of breaches will happen through poorly managed AI agent systems. Organizations that modernize identity now, Gartner says, can turn this challenge into a competitive edge.

Resilience over prevention

Cyberattacks are now seen as a normal part of doing business, not a sign of failure. Gartner analyst Leigh McMullen argued that chasing "prevention" as the only measure of success is unrealistic in the AI era. Instead, organizations should focus on resilience — limiting damage, keeping critical operations running, and recovering fast. This requires setting clear thresholds for acceptable disruption tied to what matters most to the business.

Innovation is already happening

Gartner says innovation doesn't need big new projects; it's already happening through daily fixes, workarounds, and improvements that go unnoticed. With AI-assisted tools, teams can now test defenses, simulate attacks, and refine responses without major planning cycles. By 2028, AI-driven security operations centers could cut human-handled incidents by 30%, shifting analysts from "responders" to "supervisors." Leaders, Gartner adds, must protect time for this experimentation to pay off. ■

LTM Launches BlueVerse RightLogic to tackle AI-era cyber risk

LTM launches BlueVerse RightLogic, a framework that helps enterprises assess cyber risk and securely scale AI adoption.

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

LTM, a business creativity partner for major global enterprises, has launched BlueVerse RightLogic — a new cybersecurity framework designed to help companies identify, assess, and remediate security gaps as they adopt AI at scale.

Why this matters now

AI systems can now find and exploit security weaknesses on their own. At the same time, companies face growing risks across their networks, apps, and supply chains. This has turned cybersecurity from a tech issue into a boardroom concern, as firms struggle to keep up with fast-moving threats.

"AI-driven threats are constant and scalable, requiring a new kind of defense."

— Krishnan Iyer, Chief Growth Officer, LTM

What the framework does

BlueVerse RightLogic gives businesses one clear, unified picture of their risk exposure. Instead of one-time security checks, it offers continuous, evidence-based monitoring. It looks at supply chain risks, outdated systems, network weak points, identity controls, AI-specific dangers, and how ready a company's governance is for AI risks. The model combines an "outside-in" view of how attackers see the business with an "inside-out" review of people, processes, and technology.

Krishnan Iyer, Chief Growth Officer at LTM, said AI-driven threats are constant and scalable, requiring a new kind of defense. He added that the framework helps companies measure risk, act on priorities, and build stronger security while adopting AI safely.

How it works

The program runs over 4 to 6 weeks. It includes quick diagnostics, detailed assessments, and a board-ready risk report, followed by a clear action plan. Partners then help fix issues across systems, apps, and open-source tools — making responses faster and AI adoption safer. ■

Barracuda flags threats to weak access controls

Barracuda's June 2026 report warns weak access controls and exposed remote services fuel cryptomining, botnets, and password-spraying attacks.

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

Barracuda's latest SOC Threat Radar for June 2026 highlights how attackers continue to exploit weak credentials and exposed remote access points, with three major threats identified by its Managed XDR team over the past month.

Cryptomining malware spreads via weak systems

The LemonDuck malware family was found to hijack endpoints to mine cryptocurrency and open the door to further attacks. It typically targets unpatched systems, weak passwords, and exposed services like RDP. Once inside, it uses PowerShell scripts and scheduled tasks to maintain long-term control. Barracuda recommends regular patching, limiting PowerShell access, enforcing MFA, and monitoring unusual outbound traffic.

Botnet brute-forces remote desktop access

A proactive threat hunt uncovered GoldBrute, a Java-based botnet that brute-forces RDP credentials. Once it compromises a device, that machine joins the botnet and helps scan new victims. Investigators have linked GoldBrute operators to ransomware activity, raising concerns that these infections could serve as a steppingstone to larger attacks. Businesses are urged to avoid exposing RDP directly to the internet and instead use VPNs or zero-trust access.

Password spraying from Iran targets VPNs

Barracuda also recorded a 55% jump in password-spraying attacks from Iran targeting FortiGate VPNs in May, compared with the previous month. While these attempts failed, they underline continued interest in remote access infrastructure as a target. Enabling MFA, enforcing strong password policies, and limiting login attempts remain key defenses.

Bottom line

All three threats share a common thread: weak access controls and exposed remote services remain prime targets. Barracuda stresses that visibility, MFA, and proactive monitoring are essential to staying ahead of these evolving risks. ■

The Android threat landscape expanded significantly in 2025, with fake apps and new malware families exploiting the trust users' place in official app stores.

Android malware is bypassing Google Play safeguards

Quick Heal warns that Android malware is bypassing Play Store checks by using delayed activation to steal user data and credentials.

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

Mobile malware is getting smarter, and increasingly, it's hiding in plain sight inside trusted app stores. Quick Heal Technologies has flagged a worrying trend: Android threats are now bypassing Google Play Store's security checks, putting Indian users at risk through fake apps and hidden malicious code.

What's changing?

According to Quick Heal's India Cyber Threat Report 2026, the Android threat landscape expanded significantly in 2025, with fake apps and new malware families exploiting users' trust in official app stores.

The "staged delivery" trick

The most alarming tactic is delayed activation. Apps look harmless when first submitted but can later become malicious through WebView redirects, external downloads, or location-based triggers. Some malware hides behind finance, utility, or service-themed apps before redirecting victims to fraudulent APKs designed to steal credentials.

Why it matters?

This isn't just about annoying ads or a slow phone anymore. Compromised apps can steal SMS messages, contacts, banking details, device data, and location — with some campaigns escalating to blackmail or data resale.

Staying protected

Quick Heal recommends its AntiFraud.AI tool as an added layer of defense, flagging suspicious calls, links, and app behavior in real time. More broadly, users are urged not to trust an app just because it's on a legitimate store. Basic hygiene still matters to keep devices updated, avoid sideloading apps from unknown sources, and review permissions carefully — especially requests for contacts, SMS, or accessibility access. ■

Who Owns Your Data When AI Starts Moving It?

DPDP is forcing CISOs to confront privacy risks spreading across the enterprise

By **R. Giridhar & Jagrati Rakheja** | editor.tech@timesgroup.com

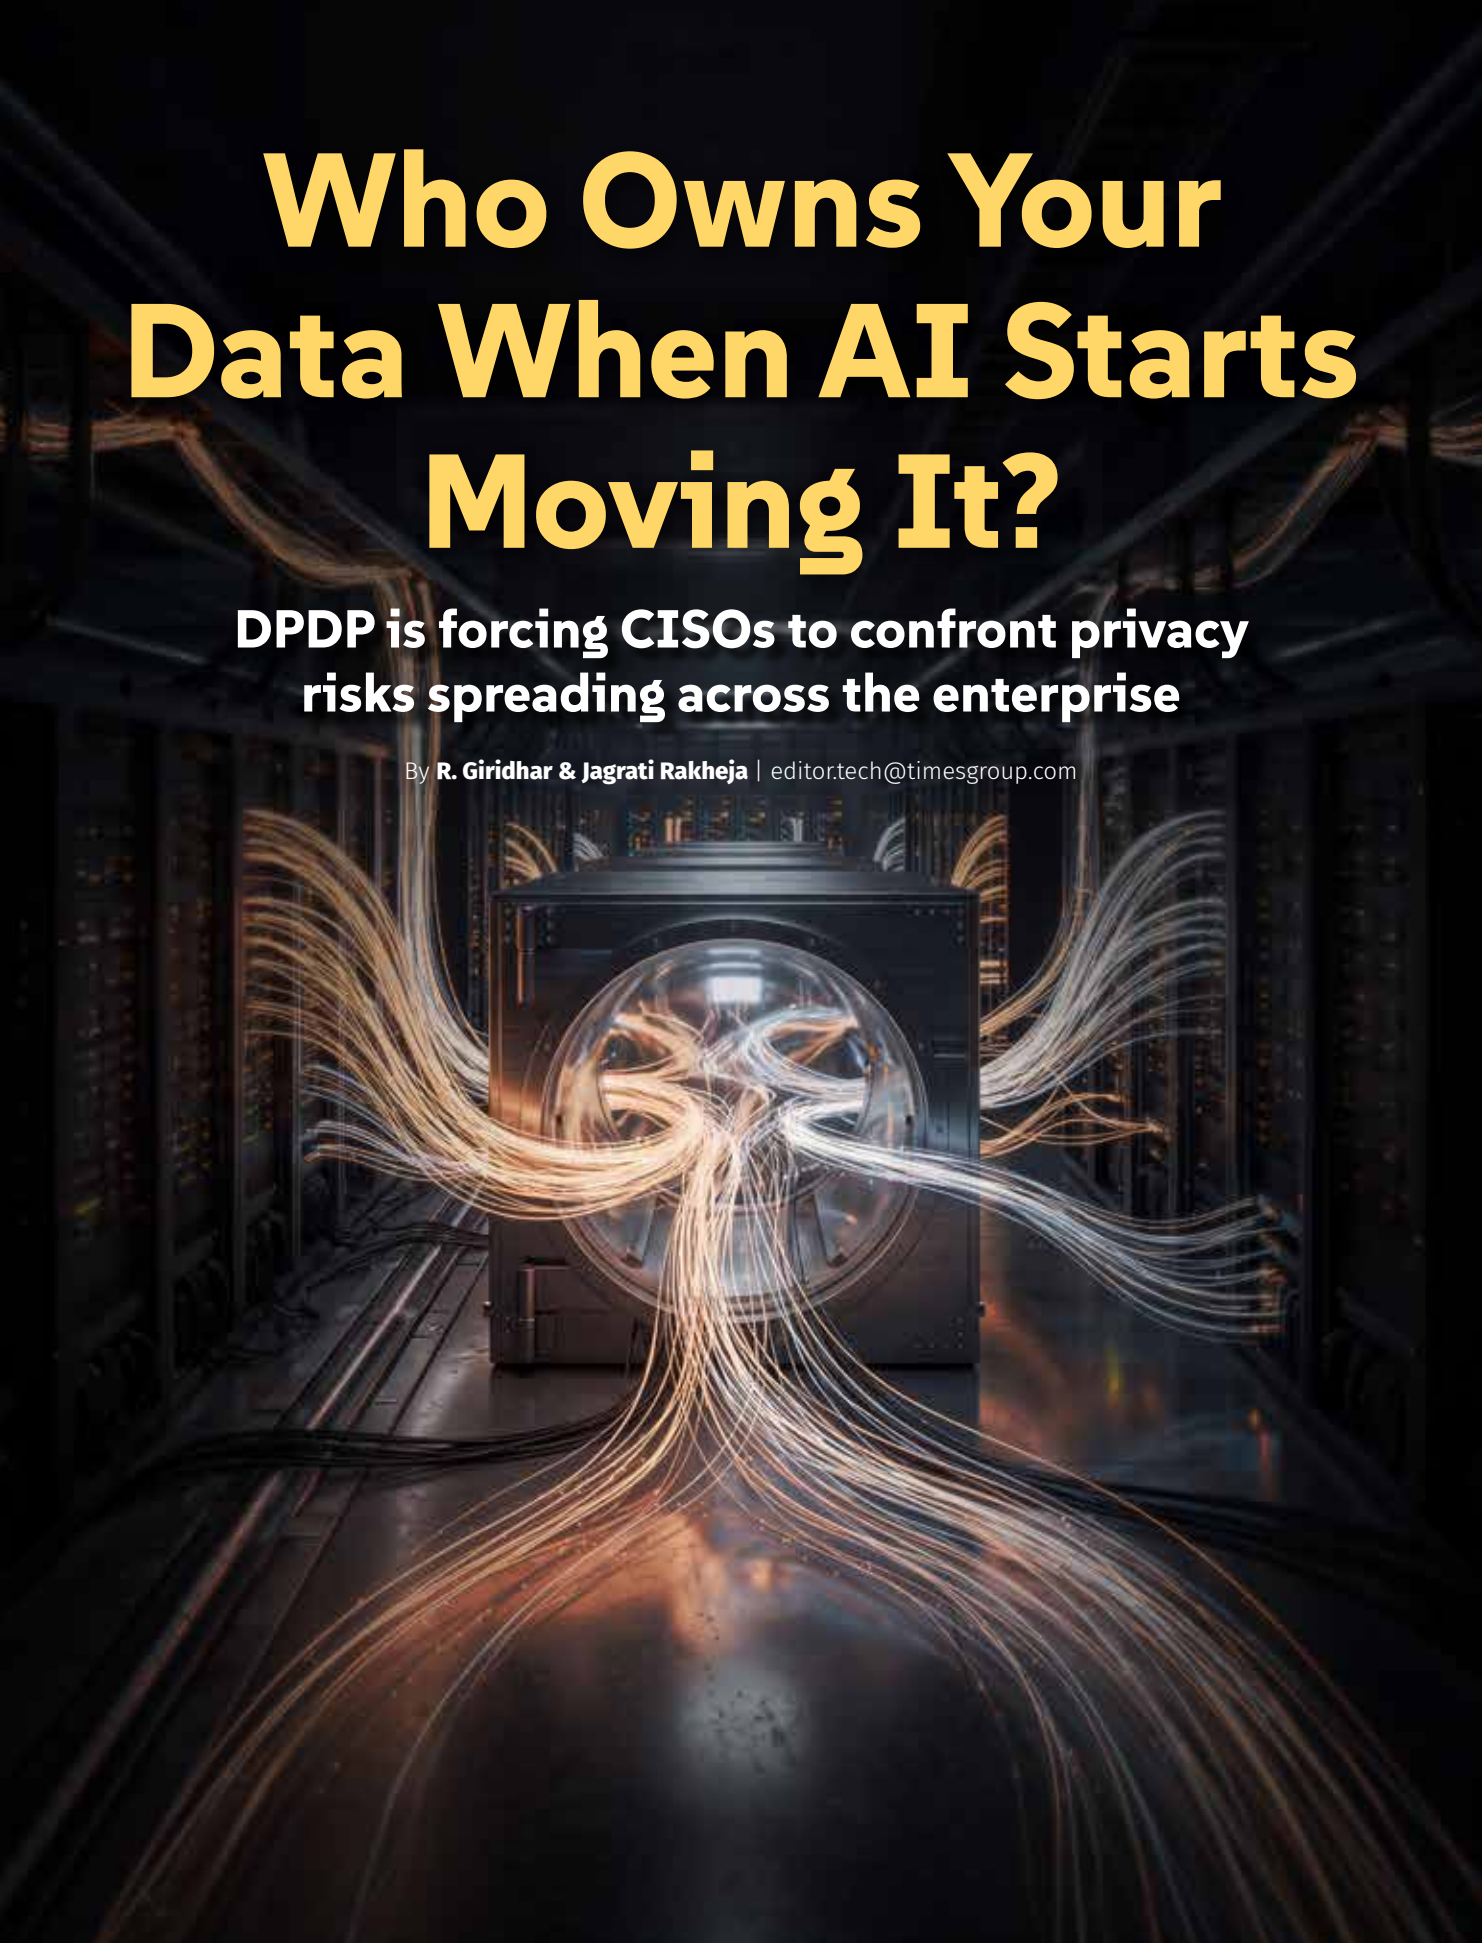

In November 2026, India's DPDP rollout reaches its first major operational milestone: the provisions governing Consent Managers come into force. The broader substantive obligations are scheduled to take effect in May 2027.

For most Indian enterprises, that may still feel like a regulatory footnote somewhere between the quarterly board meeting and the next audit cycle.

For one CISO, however, it has become a daily crisis. Call him Rajesh Mehta. The name is changed, but the scenario is a composite drawn from conversations with CISOs across banking, manufacturing, and healthcare.

Six months before the November milestone, Mehta ran a privacy gap assessment on his organisation's legacy CRM. What he found was not a compliance gap. It was a governance collapse.

Customer consent records were scattered across three systems: a Salesforce instance from 2019, a custom loan-origination platform from 2017 and a mobile-app backend launched in 2022. None could produce an auditable consent record for a specific processing activity. A customer who had consented to loan processing in 2021 had, in the system's eyes, also consented to cross-selling, affiliate data sharing, and analytics profiling.

The consent architecture had been designed for marketing convenience, not privacy governance.

The data lake was worse. Five years of Aadhaar-linked personal records sat in cold storage with no retention logic, deletion workflow or documented lawful basis for continued processing. The IT team treated the data as an archive. Under the DPDP framework, it was a liability waiting to surface.

Then came the AI problem.

DLP logs showed that three employees had pasted customer PII into a public LLM to draft personalised loan offers. The organisation's acceptable-use policy predated generative AI by four years. There was no technical enforcement at AI endpoints.

The vendor audit completed the picture. A cloud payroll provider for processing employee PII for 1,200 staff had no DPDP-specific privacy addendum. The contract, signed in 2019, contained no meaningful audit rights, sub-processor governance or data-return and deletion provisions.

Mehta's problem is India's problem.

India has spent a decade building some of the world's most ambitious digital infrastructure while often treating privacy as something that could be addressed later, by Legal, by IT, or by someone else.

The DPDP Rules, notified in November 2025, have ended that deferral.

With the first major DPDP provisions coming into force in November 2026, the broader obligations following in May 2027, penalties reaching Rs 250 crore, and privacy accountability extending across the data-processing chain, privacy is no longer simply a legal or compliance issue.

It is becoming a CISO problem. And Indian enterprises are not ready.

The Hard Part

An EY survey of more than 150 professionals, titled India's digital privacy crossroads: Understanding the DPDP Act and Rules impact and enterprise readiness, found that close to 70% were not very familiar with the DPDP Act and Rules, while more than 83% had not begun comprehensive implementation of the Act's requirements.

A parallel survey of 100 senior Indian enterprise leaders, predominantly CISOs, CIOs and CTOs, conducted by ET Edge CISO Forum and analysed by Kyndryl in July 2026, paints an equally stark picture. Only 9% describe their organisations as

audit ready. Fifty-six percent remain somewhere in the unfinished middle: 34% are still conducting initial assessments, while 22% have defined a roadmap but are yet to execute it.

The two nearest deadlines are also the weakest readiness areas.

Breach notification is the weakest dimension: 59% have no formal capability or rely on manual processes not yet aligned to the 72-hour requirement. Consent management is next, with 56% still at the awareness-or-planning stage and only 21% having reached an advanced state.

The global picture offers little comfort either. ISACA's 2026 State of Privacy report, based on more than 1,800 privacy professionals, finds that the median privacy-team size has fallen from eight to five people, even as regulatory and technological demands increase.

Cisco's 2026 Data Privacy Benchmark Study, covering more than 5,200 professionals across 12 markets, finds that 90% of organisations have expanded their privacy programmes because of AI, while 43% increased privacy spending over the past year.

And AI is only the beginning.

Agentic AI — systems capable of planning, deciding and acting across workflows with limited human intervention — will turn today's privacy exposure into something more autonomous and continuous.

The Real Question

BCG identifies five major risks in this environment: Propagation, where data moves beyond intended boundaries; Persistence, where sensitive information remains in prompts, embeddings, caches or logs; Autonomy, where agents act beyond their mandate; Emergence, where multiple agents produce unexpected outcomes; and Third-party, were external systems diffuse accountability.

Mehta's vendor problem is already Propagation and Third-party risk. As vendors introduce their own agents,

"Granular consent, managing consent revocation, and enforcing data erasure downstream—across core databases, co-lending partners, credit bureaus, NBFC networks, and third-party DSA/fintech platforms—presents immense technical friction."

Prakash Kumar Ranjan
CISO & DPO
Sammaan Capital

that can quickly become Autonomy at risk.

The real question, therefore, is not whether enterprises will comply with DPDP.

It is whether they can build the governance, technology, and organisational discipline required to become genuinely privacy capable.

What Changes Under DPDP

Indian CISOs with GDPR, CCPA/CPRA, HIPAA, PCI-DSS, ISO 27701 or NIST Privacy Framework experience are not starting from zero. Data inventories, privacy assessments, breach-response processes, processor governance and privacy-management systems all provide useful foundations.

"Our privacy journey started a few years back as we, being a global organisation, have to be compliant with GDPR and other regulations in the US and worldwide. Since the time DPDPA was announced and passed in Parliament, we had been working with the draft rules from day one to identify the possible gaps. Post the promulgation of the rules in November 2025, we have assessed our readiness, and all minor issues are being remediated. I don't see |any major roadblocks for us to be fully DPDPA compliant," says Harsha Vardhan Jonnalagadda, CISO & DPO, Birlasoft

But DPDP requires redesign in several critical areas.

The biggest is consent architecture.

GDPR provides multiple lawful bases, including Legitimate Interests. DPDP is much more consent-centric, meaning organisations that historically relied on legitimate-interest arguments for analytics, fraud detection or other activities may need to rethink those processing flows.

Consent also becomes a technical problem. Organisations need specific, documented, revocable, and auditable consent for processing activities, alongside appropriate notices and records.

The November 2026 milestone therefore cannot be treated as a policy exercise. It requires integration across CRM systems, mobile applications, branch networks, customer journeys, and third-party touchpoints.

Cross-border data flows create another layer of uncertainty. DPDP takes a negative-list approach, allowing the government to restrict transfers to countries or categories of data. That creates a strategic variable for cloud architecture, data-centre decisions, and vendor contracts.

Cisco's research shows the wider challenge: 85% of global

organisations say data localisation adds cost, complexity and risk to cross-border services.

Then there is the erasure-versus-retention problem.

DPDP gives data principals rights over their personal data, but sectoral regulators such as RBI, SEBI, IRDAI and PFRDA impose retention obligations. The result can be a direct conflict: a customer asks for deletion while another regulator requires the organisation to retain the information.

Technical control cannot resolve that contradiction.

It requires legal interpretation, a documented position, and executive approval.

The same applies to privacy ownership. The GDPR provides a more explicit framework for the independence and position of the DPO. DPDP is less prescriptive, leaving Indian enterprises to determine whether privacy sits with the CISO, Legal, a standalone DPO or a hybrid structure.

That ambiguity matters because privacy programmes can quickly lose momentum when responsibility becomes everyone's job and therefore nobody's accountability.

The Four Problems CISOs Must Solve

For CISOs, the privacy challenge breaks down into four priorities: act now, build now, resolve structurally, and monitor what comes next.

1. Act Now: Consent, Legacy Data and Vendors

Consent is the foundation.

The Kyndryl survey finds 56% of organisations remain at the awareness-or-planning stage, while only 21% are advanced.

The problem becomes even more pronounced among organisations unsure of their Significant Data Fiduciary status: 73% of this group is lagging on consent management.

That uncertainty itself is a governance risk. An enterprise cannot accurately scope its privacy

"We have a separate DPO since CISO is under the vertical of Risk function and the second line of defence."

Shermugadurai C
VP & CISO, Tamilnad Mercantile Bank

programme if it does not know which obligation tier applies to it.

The challenge is also linguistic and cultural. India has 22 scheduled languages and enormous variation in literacy, digital behaviour and data-sharing practices. For a regional NBFC or healthcare provider, producing a technically compliant notice in English is not necessarily equivalent to producing meaningful consent.

The second immediate problem is legacy data.

Years of customer information reside in backups, archives, data lakes and shadow databases. Organisations frequently cannot answer three basic questions: What personal data do we hold? Where is it? Why are we keeping it?

DPDP turns these from housekeeping questions into governance questions. For sectors built on decades-old core systems, that governance question becomes a technical one very quickly.

"The BFSI sector is generally ahead of many industries regarding baseline data governance due to existing RBI/IRDAI/SEBI regulatory frameworks. Most institutions have completed initial data mapping, inventorying, and policy alignments. The sector is now actively transitioning into technical implementation — focusing on consent management architecture, legacy data minimization, and aligning Data Fiduciary responsibilities with third-party ecosystems.

The single biggest blocker is legacy core systems and complex data lineage across interconnected ecosystems. BFSI architectures rely heavily on core banking, insurance administration, and legacy backend platforms that were built decades before privacy-by-design became standard. Mapping granular consent, managing consent revocation, and enforcing data erasure downstream — across core databases, co-lending partners, credit bureaus, NBFC networks, and third-party DSA/fintech platforms — presents immense technical friction. Upgrading these deeply intertwined systems to support real-time consent state propagation without disrupting critical, high-throughput financial operations remains the primary bottleneck for CISOs," according to Prakash Kumar Ranjan, CISO & DPO, Sammaan Capital

Cisco's research shows that although 66% of organisations have formal data-tagging systems, only 51% describe those systems as comprehensive or automated.

For Indian enterprises, legacy remediation therefore needs to go beyond deletion. It requires classification, lineage, retention logic and automated workflows.

This becomes even more important as AI systems consume enterprise data. Poorly classified historical data can become fuel for AI systems that retain, replicate, or propagate sensitive information beyond its original context.

The third immediate issue is the vendor accountability chain.

The privacy perimeter no longer ends at the enterprise firewall. It extends to cloud providers, payroll platforms, marketing systems, customer-support applications, and AI vendors.

For Indian enterprises, vendor governance is therefore not simply a procurement issue. Contracts signed years ago must be revisited — not merely for conventional data processing, but for sub-processors, public AI tools and the autonomy boundaries of future agents.

2. Build Now: The 72-Hour Clock and Shadow AI

If consent is the most immediate infrastructure challenge, breach of notification is the most visible operational test.

The DPDP Rules require a Data Fiduciary to notify the Data Protection Board of a personal data breach without delay and provide detailed information within 72 hours, unless a longer period is allowed.

That distinction matters.

Most SOCs are designed to answer: Have we been attacked?

Privacy governance requires them to answer a second question almost immediately: Does this incident involve personal data in a way that triggers privacy obligations?

In typical high-severity alerts across the sector — such as a compromised internal credential or unauthorized database query — a well-equipped BFSI SOC generally takes 4 to 8 hours to establish whether PII/NPI was exposed, exfiltrated, or merely accessed during normal operational flows.

The strain in BFSI environments rarely stems from incident detection; it centers on rapid data context identification and regulatory reporting overlaps. SIEM tools track IP traffic and access anomalies quickly but correlating an impacted server or file share back to an exact, up-to-date data classification tag — for instance, distinguishing masked account numbers from cleartext customer identifiers — requires manually cross-referencing disparate data catalogs. On top of that, BFSI security teams face the challenge of reconciling DPDP reporting expectations with existing mandates, such as CERT-In's 6-hour reporting window and sector-specific regulator notifications. "Defining whether an incident meets the legal threshold of a 'personal data breach' under DPDP — versus a standard operational security event — creates significant triage pressure between SOC teams, legal counsel, and the Data Protection Officer under tight timelines," remarks — Prakash Kumar Ranjan, CISO & DPO, Sammaan Capital.

"We know what personal data we hold, why we hold it, where it moves, who is accountable for it and how quickly we can act."

Tarovar Verma
Head of Enterprise Security, PhysicsWallah

Not every organisation experience that overlaps as friction. Some argue that anchoring to the tightest deadline on the books simplifies the rest.

"As an IT services company, RBI, SEBI, or IRDAI timelines are not applicable. However, CERT-IN, DPDPA, GDPR, CCPA, and POPIA are applicable. Since CERT-In has the least time limit, by adhering to it, others have become simple." — Vishwas Pitre, CISO & DPO, Zensar Technologies

The Kyndryl survey finds that 59% of organisations have no formal breach of capability or rely on manual processes that are not aligned to the 72-hour requirement.

The problem is particularly acute among organisations with 2,001–10,000 employees, where 76% are lagging on breach of notification readiness.

For BFSI organisations, the challenge is greater still. DPDP timelines sit alongside RBI, SEBI and IRDAI reporting requirements, potentially requiring multiple notification of workflows for the same incident.

Some organisations have not yet had to run that workflow under real conditions.

"We have not had any incident with PII data being breached. We have a robust incident response process in place to deal with incidents," declares Harsha Vardhan Jonnalagadda, CISO & DPO, BirlaSoft

Others have tested it but stop short of calling it proven.

"We have tested components of the workflow, but I would not yet describe it as fully proven end-to-end. Detection and security escalation are relatively mature; the weaker links are legal classification, rapid data scoping, processor coordination, drafting the Board submission and communicating consistently with affected data principals.

A tabletop exercise is useful, but it can create false confidence if participants are handed out clean facts. A credible test should begin with incomplete and conflicting information, involve an unresponsive processor, require the organisation to identify affected individuals from production records and make executives decide what can responsibly be communicated at each stage. Until that has been demonstrated against the clock,

the capability still exists partly on paper," says Tarovar Verma, Head of Enterprise Security, PhysicsWallah

The SOC therefore needs a privacy-classification layer — combining security, legal and business context — rather than simply adding another compliance checklist.

AI may actually help close this gap. ISACA finds that only 13% of privacy professionals currently use AI in their privacy function, although 38% plan to within 12 months. Automated privacy classification could become an important early use case.

Then there is Shadow AI.

Employees in HR, sales, finance, legal, and customer support are already using public AI tools with enterprise information. The problem is not simply that employees are breaking policy. In many cases, there is no policy designed for the technology they are using.

An employee pasting customer information into a public LLM can create a personal data transfer to an uncontracted processor.

And the same thing can happen at a vendor.

The governance response is changing. Organisations are moving away from blanket bans and towards technical safeguards, user awareness and controls at the point of interaction.

That shift is necessary because agentic AI will make the problem continuous. A human entering a prompt is one thing. An agent with persistent access to enterprise data is another.

The CISO needs to know not just what employees are entering into AI systems, but what data AI systems can access, retain, retrieve and act upon.

3. Resolve Structurally: Regulation and Ownership

The hardest privacy problems cannot be solved by buying another tool.

The multi-regulator conflict is one.

DPDP may require erasure while sectoral regulation requires retention. Privacy rules may require granular consent while legacy sector practices rely on broad consent. Cross-border requirements may intersect with existing localisation obligations.

"Depending on the size of the organisation, you may have a different structure in each case, or even the CISO and DPO can be the same person carrying out different responsibilities."

Vishwas Pitre
CISO & DPO,
Zensar Technologies

The answer cannot be another spreadsheet mapping one regulation to another.

CISOs need a unified data-risk taxonomy that brings DPDP, sectoral requirements, and AI governance into a common control framework.

BCG's agent AI research reinforces the need. As agents cross organisational and functional boundaries, separate privacy, security, data and AI governance silos become increasingly difficult to sustain.

The second structural problem is ownership.

Data from the Kyndryl survey offers important insights. Organisations with a dedicated DPO/CPO role are the strongest performers: 50% are advanced, compared with 34% where privacy sits with the CISO and 27% under joint cross-functional ownership.

Most strikingly, organisations with no designated privacy owners are 80% lagging, with zero advanced.

The lesson is straightforward: the precise structure matters less than having clear accountability.

Three models are emerging.

Privacy can sit under the CISO, bringing technical coherence but risking the subordination of privacy rights to security priorities.

It can sit under Legal, strengthening regulatory interpretation but potentially disconnecting privacy governance from the systems where data resides.

Or it can be a standalone, board-reporting function, providing independence but risking becoming a figurehead without operational authority.

In practice, Indian enterprises are already testing each of these models.

"We have a separate DPO since the CISO is under the vertical of Risk function and the second line of defence," Shermugadurai C, VP & CISO, Tamilnad Mercantile Bank

"The responsibilities of the CISO and DPO are different, and hence there is no conflict. Depending on the size of the organization, you may have a different structure in each case, or even the CISO and DPO can be the same person carrying out different responsibilities," clarifies Vishwas Pitre, CISO & DPO, Zensar Technologies

For many large enterprises, the most practical model may be hybrid: a privacy function with clear accountability, co-ordination between CISO and General Counsel, embedded technical expertise, and a dotted-line relationship with the board.

The critical ingredient is executive sponsorship.

ISACA finds that while 56% of

boards say privacy is adequately prioritised, only 22% expect privacy budgets to increase, while half anticipate decreases.

Board recognition without budget is not governance.

4. Monitor: Cross-Border and Public Infrastructure Risk

Not every issue requires immediate restructuring.

Cross-border data flows are one example.

DPDP's negative-list approach creates uncertainty, but until specific restrictions emerge, organisations with significant offshore dependencies should focus on documentation and contingency planning.

The strategic question is larger than India. Cisco finds that 83% of organisations believe harmonised global data protection standards would reduce the need for country-specific localisation laws.

Indian enterprises should therefore build privacy architectures that can adapt rather than locking themselves into one regulatory scenario.

Public digital infrastructure creates another inherited risk.

Enterprises increasingly rely on Aadhaar, DigiLocker, Ayushman Bharat, and other digital public infrastructure. These systems can create vulnerabilities outside an enterprise's direct control.

The enterprise cannot necessarily remediate the underlying infrastructure, but it can acknowledge the risk, strengthen controls at its own boundary, monitor anomalous activity and ensure its contracts and incident-response plans account for external dependencies.

The Hardest Part: Who Owns Privacy?

Technology is not the hardest privacy problem.

Ownership is.

Consent architecture can be engineered. Legacy data can be classified and deleted. Vendor contracts can be renegotiated. SOC workflows can be redesigned.

"We know what personal data we hold, why we hold it, where it moves, who is accountable for it and how quickly we can act."

Harsha Vardhan Jonnalagadda
CISO & DPO, Birlasoft

What cannot be automated is organisational alignment.

A CISO who approaches the board asking for money for "DPDP compliance" may receive a compliance budget.

A CISO who explains that the organisation has unquantified vendor liability, uncontrolled employee use of AI, five years of unclassified personal data and no defensible answer to a regulator asking where customer information went is asking for something different: a governance mandate.

That distinction is crucial.

Privacy is increasingly becoming a shared responsibility across Security, IT, Legal, HR, Marketing, Product and Operations. But shared responsibility without a clear owner creates paralysis.

The CISO has a strong claim to privacy ownership because much of the work is technical: encryption, access control, data discovery, consent infrastructure, breach detection and AI endpoint controls.

Yet Legal must own regulatory interpretation. Business functions own how data is used. Data teams' own classification and lineage.

The answer is therefore not to force privacy into a single silo, but to establish clear accountability across the ecosystem.

For Significant Data Fiduciaries in particular, a hybrid model makes sense: a dedicated privacy function, technical privacy expertise embedded in IT and Security, strong Legal involvement, and direct board visibility.

What matters most is that someone can answer the question:

Who is accountable when personal data moves somewhere it should not?

If the answer is unclear, the enterprise is not ready for privacy.

"The real worst case is the absence of a DPDP programme that appears complete — until the first real incident happens. 13 May 2027 is a reality that is driving us towards a specific goal, where obligations have become part of normal business operations. We know what personal data we hold, why we hold it, where it moves, who is accountable for it, and how quickly we can act when something goes wrong. Consent withdrawal, data-principal requests, retention and deletion, processor oversight and breach notification are measurable workflows rather than policy statements.

In that best case, ownership is also unambiguous. Legal interprets the obligation; privacy defines the control framework, security detects and contains incidents, technology operates the required systems, and business owners remain accountable for the processing they sponsor. Senior management receives evidence of readiness — test results, response times, control exceptions

and remediation status — not simply a compliance percentage.

In modern organisations where a lot of moving parts are present, where business is not keeping track of data usage, security controls are not configured to tap into business changes, where consent inventories are not well kept — Legal and all concerned stakeholders believe the DPDP implementation group has completed the job and we are compliant, only to wake up to a rude incident," Tarovar Verma, Head of Enterprise Security, PhysicsWallah.

The Privacy Opportunity

There is a temptation to frame DPDP as another compliance burden.

That would be a mistake.

Done properly; privacy capability is fundamentally disciplined.

An enterprise that builds consent architecture also builds better data lineage. One that remediates legacy data reduces its attack surface. One that knows where personal data resides can respond faster to incidents. And one that governs AI endpoints is better positioned for agentic AI.

Privacy and security are therefore not separate disciplines. Privacy governance is, in many ways, security governance with personal data as its primary object.

The business case is also becoming clearer.

Cisco's 2026 benchmark finds that 99% of organisations report at least one measurable benefit from privacy investments. Ninety-five percent report stronger customer trust, while 96% say enhanced data controls have unlocked agility and innovation.

That changes the board conversation.

Privacy is not simply about avoiding a Rs 250-crore penalty. It is about building an enterprise capable of knowing what data it has, why it has it, where it moves, who can access it, and what AI systems can do with it.

That capability becomes increasingly valuable as enterprises move from copilots to autonomous agents.

The organisations that build privacy discipline now will enter the agentic era with cleaner data, better controls, faster incident response and stronger customer trust.

The ones that postpone it will eventually build the same capabilities — but at higher cost, under regulatory pressure and after a preventable incident has forced the issue.

The privacy reckoning, then, is not really about DPDP.

It is about whether Indian enterprises are prepared to treat personal data as an asset that requires stewardship rather than raw material that can move freely through increasingly autonomous systems.

The question is no longer whether enterprises will build privacy capability.

It is which CISOs will build it first — and what they will build between now and May 2027 that the rest of the market will spend years trying to replicate.

The window to May 2027 may look like 18 months. For enterprises with fragmented data, legacy systems, hundreds of vendors, and rapidly expanding AI use, it is considerably shorter. ■

CISO Action List

Priority	What CISOs should do
1. Make November 2026 a hard milestone	Treat Consent Manager for readiness as an engineering and governance programme, not a policy exercise.
2. Know your data before you govern it	Prioritise personal-data discovery, classification, lineage, retention and deletion across legacy environments.
3. Rebuild the vendor chain	Review critical contracts for DPDP accountability, sub-processors, AI usage, audit rights, and deletion obligations.
4. Put privacy inside the SOC	Build the capability to identify personal data incidents rapidly and operate against the 72-hour reporting requirement.
5. Map Shadow AI now	Know which AI tools employees and vendors are using, what personal data they process, and what technical controls exist at the endpoint.
6. Resolve regulatory conflicts at board level	Document positions on retention versus erasure, cross-border transfers and sector-specific obligations with Legal and executive sign-off.
7. Establish clear privacy ownership	A dedicated privacy function with strong CISO, Legal and board engagement is preferable to a model where everyone owns privacy, and nobody is accountable.
8. Reframe privacy as an enterprise capability	The strongest argument for investment is not compliance. It is cleaner for data, lower risk, faster response, responsible AI, and customer trust.

When resilience becomes architecture: India's cybersecurity reckoning

Sunil Sharma, Managing Director & Vice President – Sales, India & SAARC, Sophos, delivers a frank diagnostic on ransomware, identity security, and why resilience is an architectural, not reactive decision.

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

India's ransomware numbers tell a story of cautious progress median payments are down sharply, and more organizations are holding the line. But Sunil Sharma, Managing Director & Vice President – Sales, India & SAARC at Sophos, isn't ready to call it a turning point. Not yet.

In a wide-ranging conversation with CISO Forum, Sharma cuts through the optimism with frontline data and uncomfortable truths: 53% of Indian victims still paid the ransom in 2025, identity breaches are accelerating faster than governance frameworks can respond, and the CISO role itself is structurally broken at a 10,000:1 ratio to the businesses it serves.

What emerges is less a vendor pitch than a frank diagnostic — on where Indian enterprises are genuinely maturing, where dangerous assumptions persist, and why resilience, at its core, is an architectural decision that no amount of reactive spending can substitute for.

CISO FORUM: Your own State of Ransomware 2025 report found that 53% of Indian victims paid the ransom, yet median payments dropped sharply. What does this signal about where Indian enterprises stand on the prevention-to-resilience continuum — and what is you seeing on the ground that the numbers don't fully capture?

SUNIL SHARMA: The data tells a story of transition, not triumph. The median ransom payment in India dropped by 79% to US$481,636, and the proportion of organizations paying the ransom fell from 65% in 2024 to 53% in 2025, a meaningful directional shift. But what the numbers don't fully capture is the nature of that payment decision. Indian businesses spent an average of US $1.01 million on recovery costs, excluding the ransom itself, which means paying the ransom is often the beginning of the financial pain, not the end of it.

On the ground, the market is split: a small but growing; cohort has invested in detection, backup maturity, and MDR is genuinely more resilient, and their outcomes are measurably better. But the majority are still in reactive mode; they pay because it seems like the fastest path back to operations. From an operational perspective, 41% of organizations from our study that experienced a ransomware attack cited a lack of people or capacity as a root cause of their attack, and 39% acknowledged that not having the right cybersecurity products contributed to their being victimized.

Most organizations that pay the ransom do so not because it is the smartest business decision, but because inadequate prior investment in prevention, backups, and response readiness left them with no other viable option. The most dangerous thing about India's ransomware picture right now is the gap between awareness and action. People know what good looks like, but internal capacity constraints mean the journey from knowing to doing is slower than attackers are moving.

CISO FORUM: Sophos's 2026 Active Adversary Report finds 67% of all incidents are identity-rooted, with attackers reaching Active Directory in under 3.5 hours. Yet, our survey shows that identity lifecycle governance is among the lowest-maturity capabilities among Indian CISOs. What is the most dangerous identity assumption that Indian enterprises are still operating with?

SUNIL SHARMA: The most dangerous identity assumption Indian enterprises are still operating with is that Active Directory governance is an IT operations problem, not a security priority.

According to Sophos's newly released State of Identity Security 2026 report:

- 76.8% of organizations surveyed in India suffered at least one identity-related breach in the past year
- 79% of ransomware victims in India confirmed their ransomware incident stemmed directly from an identity attack, establishing identity compromise not just as an entry vector, but as the primary delivery mechanism for the most damaging threat Indian enterprises face

The 2026 Active Adversary Report reinforces why this is so consequential. Globally, 67% of all incidents investigated by our IR and MDR teams were rooted in identity-related attacks, with attackers exploiting compromised credentials, weak or missing MFA, and poorly protected identity systems, often without deploying any new tools or techniques. Once inside, attackers reach Active Directory in a median of 3.4 hours, and the median dwell time has declined to three days, meaning the window for human-driven response is closing faster than most teams are prepared for.

The Two Most Dangerous Assumptions:

- **Deploying MFA equals identity security:** Adversary-in-the-middle phishing kits are now commodity tooling. They proxy a counterfeit login page, capture both the credentials and MFA token in real time and replay them to the legitimate identity provider within seconds. App-based push and SMS one-time codes do not stop them.
- **Credentials not in a breach database are safe:** Sophos' telemetry shows stolen credentials on the dark web more than doubled in a single year. If credentials have not appeared in a known breach database, that does not mean they are safe.

CISO FORUM: Talent shortage has surpassed budget as the #1 internal barrier in our survey — 50% of Indian CISOs rate it high or

"The most dangerous thing about India's ransomware picture right now is the gap between awareness and action."

extremely severe. Sophos MDR now secures over 33,000 organizations globally. How should a CISO frame the MDR decision to their board — not as outsourcing, but as a strategic architecture choice — and where does human-led MDR end and the enterprise's own team begin?

SUNIL SHARMA: A CISO needs to make it clear to the board that MDR is not outsourcing security. It is how modern security operations are architecturally designed when talent scarcity is a structural reality. The board already accepts that the organization does not run its own data center for email or build its own ERP from scratch. Security operations at 24/7 coverage and threat-hunting depth have the same economics: it is more effective and more sustainable when delivered through a model built for scale.

Sophos MDR and XDR are trusted by more than 75,000 organizations worldwide, and our analysts' experience securing those organizations from advanced, human-led threats directly informs AI Assistants and detection capabilities that no internal team could build unilaterally.

The Two-Layer Operating Model:

- **Sophos MDR -the operational execution layer:** Handles 24/7 monitoring, threat detection, investigation, and active response.
- **Enterprise internal security team - the strategic layer:** Owns governance, risk management, compliance, vendor relationships, policy, and the contextual business knowledge that no external team can possess.

This is not an abdication. It is how elite security is delivered at scale. The talent shortage is a permanent condition of the market, not a temporary one. The architecture should be designed around that reality.

CISO FORUM: Sophos acquired Secureworks in February 2025, bringing Taegis MDR and XDR into the fold. For an Indian enterprise CISO evaluating this combined platform, what has materially changed in detection, response, and coverage depth — and how does Sophos ensure that integration complexity doesn't create new blind spots during the transition?

SUNIL SHARMA: With the completion of the Secureworks acquisition in February 2025, Sophos became the leading pure-play cybersecurity provider of MDR services globally. For an Indian enterprise CISO evaluating the combined platform, three things have materially changed:

- **Detection depth:** Secureworks' Counter Threat Unit (CTU), which tracked more than 150 threat groups, is now part of Sophos X-Ops. The threat of intelligence engines behind every detection has expanded dramatically.
- **Platform integration:** Sophos Endpoint is now natively integrated and automatically included in all Taegis XDR and Taegis MDR subscriptions, giving customers immediate access to combined prevention, detection, and response capabilities in a single platform while lowering costs and simplifying operations.
- **Coverage breadth:** The combined portfolio now includes Managed Services, next-generation SIEM, and ITDR capabilities that previously required separate vendor relationships.

On integration complexity and blind spots - Any acquisition introduces a transition period. The mitigation is architectural clarity. Customers on either platform retain continuity of their existing detection and response workflows, with Taegis remaining an open platform and Sophos Central serving as the unified management layer. The integration is being sequenced deliberately rather than rushed, specifically to avoid the coverage gaps that rapid integration can create.

CISO FORUM: Sophos launched Identity Threat Detection and Response in October 2025. Our survey shows 68% of Indian CISOs rate identity-centric breaches as high or extremely severe, yet ITDR remains a nascent investment category in India. What should a CISO do in the next 90 days to reduce identity exposure — before a formal ITDR program is in place?

SUNIL SHARMA: Before any formal ITDR program is in place, a CISO can take three high-impact actions in 90 days that will materially reduce identity exposure:

- **Run an identity posture assessment immediately:** Without a full ITDR deployment, run a manual review of dormant accounts, service account privileges, and MFA gaps against your Active Directory and Entra ID environment. This will surface the highest-risk misconfigurations before they are exploited.
- **Check credential exposure on the dark web:** Sophos X-Ops observed a 106% increase in stolen credentials for sale on the dark web between June 2024 and June 2025. Sophos ITDR scans the dark web and breach databases for evidence of leaked or stolen credentials. Identify when login credentials are exposed on the dark web and breach of databases.
- **Enforce phishing-resistant MFA on highest-privilege accounts:** Specifically, accounts with access to Active Directory, cloud admin consoles, and financial systems. Given that attackers are reaching Active Directory in a median of 3.4 hours, hardening those accounts is the single most effective thing an organization can do before a formal ITDR program is funded and deployed.

CISO FORUM: Sophos's sector research shows manufacturing is disproportionately targeted, with exploited vulnerabilities driving 32% of attacks, while BFSI leads in attack frequency in India. Given the layered regulatory pressures on both RBI, IRDAI, and now DPDP — how does a resilience architecture for these two sectors differ, and where do most organizations get it wrong?

SUNIL SHARMA: The two sectors have fundamentally different threat profiles and therefore require different resilience architectures:

Manufacturing - According to the 2025 Sophos State of Ransomware report:

- **Root cause:** Exploited vulnerabilities are the leading root cause, responsible for 32% of incidents, while lack of expertise is the most common organizational factor, cited by 42.5% of victims.
- **Positive shift:** 50% of manufacturing organizations stopped attacks before data could be encrypted, more than double the previous year's 24%.
- **Evolving threat:** Adversaries have adapted. 39% of manufacturers that experienced encryption also had data stolen, one of the highest rates across all surveyed sectors.
- **Resilience architecture priority:** OT/IT network segmentation, aggressive vulnerability patching on edge devices and legacy systems, and backup strategies that survive a double-extortion scenario where data has already been exfiltrated before encryption begins.

BFSI - The sector experienced the highest frequency of ransomware attacks in India over the past year.

- **Regulatory complexity:** The sector faces a layered regulatory environment. RBI's cybersecurity framework, IRDAI's guidelines, and DPDP Act obligations mean the resilience architecture must simultaneously satisfy operational recovery, data integrity, and regulatory reporting timelines.
- **Where most organisations get it wrong:** Resilience is treated as a technology checklist rather than an operational capability. Two critical distinctions:
  - Having backups is not the same as having tested, recoverable backups
  - Having an IR plan is not the same as having an IR plan that has been simulated under realistic adversarial conditions

CISO FORUM: Tool proliferation and alert fatigue rank as the third-most-severe internal challenge in our survey. Sophos has deployed an AI Assistant for case investigations alongside its human analyst teams. Where is AI genuinely accelerating security outcomes in your MDR operations — and where does the over-reliance on AI actually create risk that only human judgment can address?

SUNIL SHARMA: With Sophos MDR analysts securing more than 35,000 organizations globally, there is a clear view of where AI is genuinely moving the needle and where human judgment remains irreplaceable. AI is accelerating security outcomes in three key areas:

- **Initial triage and alert enrichment:** AI processes and prioritizes incoming alerts at a speed no human team can match, ensuring analysts focus on what matters most.
- **Pattern recognition across large telemetry volumes:** AI identifies patterns across data volumes that would overwhelm human analysts working at speed.
- **Case summarisation:** AI reduces the cognitive load on analysts before they begin an investigation, helping security teams quickly identify risks, enrich investigations with threat intelligence, and take faster remediation actions.

And where over-reliance on AI creates risk:

- **Adversary behavior at decision points:** The Sophos Active Adversary Report 2026 found no evidence of a major AI-driven

"Every compliance investment should be evaluated against whether it would have reduced the blast radius of the last breach, or the next one. If the honest answer is no, it is an audit exercise, not a security investment."

transformation in attacker behavior. AI is adding scale and noise but not yet replacing attackers. The fundamentals still matter about strong identity protection, reliable telemetry, and the ability to respond quickly when something goes wrong. An AI system trained on historical attack patterns will have systematic blind spots around novel techniques.

- **Business context and judgment:** Human judgment is irreplaceable when understanding business context, knowing that a particular anomalous access event at 2 am is a legitimate financial close process or that isolating a specific endpoint will halt a critical production line. That contextual business knowledge cannot be encoded in any model.

The architecture that works is AI for speed and scale, humans for judgment and context.

CISO FORUM: 60% of Indian CISOs in our survey rate DPDP Act enforcement as having a high or extremely high impact on their 2026 planning. Sophos recently expanded DNS Protection management regions to India. How should CISOs think about the intersection of data protection compliance and security architecture — and what compliance-driven investments are also genuinely improving security posture versus those that only satisfy auditors?

SUNIL SHARMA: The DPDP Rules, notified on November 14, 2025, provide an 18-month phased compliance timeline and require data fiduciaries to issue clear, standalone consent notices explaining the specific purposes for which personal data is being collected. The intersection of DPDP compliance and security architecture is more direct than most CISOs currently appreciate.

Investments that genuinely improve security posture while satisfying compliance:

- **Data classification and mapping:** You cannot protect what you have not identified. This is both a compliance requirement and a foundational security control.
- **Breach detection and notification capabilities:** DPDP mandates prompt notification, which requires the detection of infrastructure to exist first. Compliance here directly drives security investment.
- **Access controls and encryption:** Both are regulatory requirements and foundational security controls, making them dual-purpose investments.

Investments that tend to satisfy auditors without improving posture:

- Point-in-time compliance assessments
- Checkbox policy documentation without operational implementation

- Vendor risk questionnaires that are never verified.

The DNS Protection expansion to India-based management regions is one concrete example of an architectural decision that serves both compliance intent, keeping data processing within jurisdictional boundaries, and the security function.

The right framing for the board is: every compliance investment should be evaluated against whether it would have reduced the blast radius of the last breach, or the next one. If the honest answer is no, it is an audit exercise, not a security investment.

CISO FORUM: Our survey reveals that 48% of CISOs rate inadequate IR testing as a severe challenge, and most organizations' playbooks have never been tested under realistic conditions. Sophos launched Emergency Incident Response in 2025. What does your frontline IR data tell you about where Indian enterprise response plans most predictably fail — and what should a CISO simulate before a real incident forces the lesson?

SUNIL SHARMA: From Sophos's global frontline IR and MDR data spanning 661 cases across 70 countries, response plans most predictably fail at three points. These failure patterns are consistently observed across markets, including India, where survey data confirms that 48% of CISOs rate inadequate IR testing as a severe challenge.

The three most common failure points are:

- **Decision authority:** When an incident is active, and containment requires isolating systems, shutting down services, or notifying regulators, organizations repeatedly discover that their IR plan does not clearly assign who has the authority to make those decisions. The technical response may be ready, but the governance response is not. Research consistently shows that the leading cause of prolonged incidents is not attacker sophistication; it is internal indecision. The technical problem might be resolved in hours, while the leadership problem drags on for days.
- **Log availability:** Missing logs due to data retention issues doubled in the most recent Active Adversary dataset, largely driven by firewall appliances where the default retention was only seven days and, in some cases, 24 hours. Forensic investigation is impossible without telemetry. Before simulating an incident, validate that logs are being retained and are accessible.
- **Backup integrity under adversarial conditions:** Most organizations have never tested whether their backups can be restored at speed when an adversary has had 72 to 96 hours of access to the environment. The simulation a CISO should run before a real incident is a tabletop exercise that begins with "your backups are inaccessible" and works forward from there. That single constraint will expose every assumption in your recovery plan.

CISO FORUM: Your CISO Report 2026 puts the global CISO-to-business ratio at 10,000:1 — an unsustainable structural imbalance. Our survey corroborates this: burnout and attrition in security teams are rated high severity by 39% of respondents, and talent is now the top barrier. What is your honest assessment of how Indian enterprises should restructure the security function not just hire more to make the CISO role and the teams beneath it sustainable at the speed AI-era threats demand?

SUNIL SHARMA: The current model, one CISO supported by an overextended team, accountable for an expanding surface area of risk, and operating without adequate tooling or staffing, is not a talent problem. It is a structural design problem, and hiring more people into a broken structure will not fix it.

The 2026 CISO Report shows that organizations can no longer rely on traditional security leadership models to keep pace with the scale and sophistication of today's threats. There are only 35,000 CISOs worldwide serving an estimated 359 million businesses, a 10,000:1 ratio that represents a market failure.

For Indian enterprises, the restructuring required has three components:

- **Separate strategy from operational security:** The CISO's role should be risk governance, regulatory navigation, board communication, and security architecture, not operational triage. Operational security should run through MDR or an equivalent 24/7 model that removes the CISO from the on-call rotation.
- **Invest in automation at the operational layer:** Junior analysts should not be manually doing work that degrades their capability and accelerates attrition. Burnout has measurable consequences. Nearly half of affected professionals report heightened anxiety about breaches, 39% admit to reduced productivity, and a third report reduced engagement at work, all of which directly undermine the effectiveness of the very defenses they are responsible for.
- **Redefine what "the security team" means:** In the AI era, the team includes managed services, automated detection pipelines, and the intelligence of platforms like Sophos X-Ops. The CISO who tries to build everything in-house will most likely lose to the adversary every time. The CISO who architects the right combination of internal governance, external operations, and platform intelligence has a sustainable model. ■

Left Behind by Design: India's AI Security Gap

India has entered Project Glasswing, but as a second-wave participant rather than a founding partner. While US and UK organisations gained early access to advanced cybersecurity capabilities, Indian agencies tested less capable alternatives. The gap is narrowing, but the delay matters. Indian CISOs now face challenges in access, regulation, and talent.

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

The Models: What Has Actually Changed

Claude Mythos Preview and GPT-5.5 Cyber (referred to herein as GPT-5.5 Cyber or Daybreak) are not enhanced signature engines or smarter Security Orchestration, Automation and Response (SOAR) playbooks. They are autonomous reasoning systems capable of end-to-end offensive and defensive cyber operations — and their capabilities have been independently verified at scale.

Anthropic launched Project Glasswing in April 2026, initially restricting Mythos Preview to approximately 50 partner organisations. The rationale was unambiguous: the model could autonomously discover and exploit software vulnerabilities without human guidance, and releasing it publicly risked placing a supercharged cyberweapon in adversarial hands [1]. Within four weeks, those 50 partners had collectively surfaced over ten thousand high- or critical-severity vulnerabilities across the most systemically important software in the world [1].

The numbers are operationally staggering. Cloudflare found 2,000 bugs — 400 of them high- or critical-severity — with a false-positive rate that its team rated better than human testers [1]. Mozilla fixed 271 vulnerabilities in Firefox 150, more than ten times the number found in the previous release [1]. The UK's AI Security Institute confirmed Mythos Preview is the first model to autonomously complete both of its cyber ranges — multistep simulations of full corporate network attacks — end to end [1].

OpenAI's response — GPT-5.5 Cyber, announced on 7 May 2026 — takes a different structural approach [6]. Rather than a single restricted model, OpenAI created a Trusted Access for Cyber (TAC) framework: a tiered identity-and-trust system that gates capability by the verified status of the user. Standard GPT-5.5 handles general defensive tasks. GPT-5.5 with TAC unlocks vulnerability triage, malware analysis, binary reverse engineering, and patch validation for verified defenders. GPT-5.5 Cyber — available in limited preview to critical infrastructure operators — adds authorised red-teaming and live exploit validation [6]. Phishing-resistant authentication is mandatory for the most permissive tier [6].

> ***Progress on software security used to be limited by how quickly we could find new vulnerabilities. Now it's limited by how quickly we can verify, disclose, and patch the large numbers of vulnerabilities found by AI.***
>
> ***— Anthropic, Project Glasswing: An Initial Update, May 2026 [1]***

What both programmes share is a foundational claim described by LSE researchers Dr Beatriz Lopes Buarque and Microsoft AI engineer Abdullah Abu-Hassan as a collapse in the cost of offensive cybersecurity [5]. Capabilities that once required an elite operator with decades of experience can now be executed with a few lines of code. An elite cyber capability has become, in effect, a software function. The Mythos system card confirms the model can run end-to-end attacks on enterprise networks with weak security postures, and complete a 32-step corporate network attack simulation in timeframes that would take a skilled human twenty hours [5].

India's Position: Inside the Perimeter, Behind the Curve

India is part of Project Glasswing. That much is confirmed. What is less discussed — and more consequential for enterprise CISOs — is the nature of that participation and what it reveals about India's structural position in the frontier AI security landscape.

The Access Lag

When Anthropic launched Project Glasswing in April 2026, the founding cohort of approximately 50 organisations was drawn almost entirely from the United States and United Kingdom [7]. Amazon, Google, Microsoft, Apple, NVIDIA, CrowdStrike, Palo Alto Networks, and the UK's AI Security Institute were scanning and hardening critical software with Mythos Preview while the rest of the world waited. India was not at that table.

During those weeks, Indian organisations that recognised the risk were conducting controlled security tests using Claude Opus 4.7 — Anthropic's preceding flagship model — as a substitute [4]. Major IT services firms including Infosys and Tata Consultancy Services ran these tests. CERT-In reviewed critical systems including Aadhaar and government login platforms using the same [4]. That is the access gap: India's most important cybersecurity institutions spent the most critical weeks of the frontier AI security transition working with a less capable tool, scanning the same infrastructure that a Mythos-class adversary could target.

By late May, Anthropic expanded the programme to approximately 150 organisations across more than 15 countries [7]. India was included in this second wave. According to reporting by The Indian Express, the Indian organisations now understood to have received access to Mythos Preview include: the Indian Cyber Crime Coordination Centre (I4C); CERT-In; the National Critical Information Infrastructure

Protection Centre (NCIIPC); and the Department of Telecommunications' Digital Intelligence Platform (DIP) [7]. NCIIPC — which reports to the National Security Advisor in the Prime Minister's Office — and CERT-In had specifically requested access to scan banking and power infrastructure [7]. Discussions are also underway to extend access to cybersecurity research institutions and dedicated AI security teams within India's largest IT services companies [7].

Inclusion in the second wave is meaningful. It is not equivalent to founding membership. The organisations that joined in April have six to eight weeks of operational experience with Mythos Preview — six to eight weeks of vulnerability findings surfaced, triaged, and patched — that Indian institutions do not. Every day of that lag is a day during which vulnerabilities in systems that underpin Indian enterprise environments sat unexamined by the most capable scanning tool ever deployed. Some will already be in Anthropic's coordinated disclosure queue. Others will not surface for months.

The Sovereign Hosting Constraint

India's access position is further complicated by the government's push for sovereign hosting of Claude AI — specifically for use in banking, telecom, and critical infrastructure, where jurisdictional and national security concerns around foreign-hosted AI infrastructure are acute [4]. This ambition is legitimate and directionally correct. Its near-term consequence is a meaningful delay in the timeline for regulated sector organisations to deploy Mythos-class capability at production scale within compliant data boundaries.

For CISOs in banking, insurance, and telecom — the sectors where Mythos-related regulatory guidance from SEBI and RBI is already emerging — the sovereign hosting question is not an abstract policy debate. It is the difference between being able to deploy AI-assisted vulnerability scanning on production systems and being limited to air-gapped pilots on non-production environments. Until the sovereign hosting framework is resolved, regulated CISOs face a compliance constraint that their global peers do not.

India's Glasswing Cohort: What We Know

- Government entities confirmed to have access: I4C, CERT-In, NCIIPC, DoT's Digital Intelligence Platform [7].
- Cybersecurity-focused research institutions are also understood to have received access [7].
- Discussions ongoing for access for AI and cybersecurity teams within major Indian IT services firms [7].
- Finance Minister Sitharaman and IT Minister Vaishnaw convened high-level meetings in April to assess Mythos risks to the banking sector [7].
- IBA has been directed to develop a coordinated institutional response mechanism [7].
- India is seeking sovereign hosting of the Claude AI model, citing jurisdictional and national security concerns for regulated sectors [4].
- SEBI has formed cyber-suraksha.ai, a task force to assess frontier AI risk to regulated entities [4].
- RBI is coordinating with global regulators on Mythos-related risks [4].

Sources: The Indian Express, June 2026 [7]; Inc42, June 2026 [4].

The Adversary Does Not Have the Same Constraint

The access lag and the sovereign hosting delay apply to Indian defenders. They do not apply to adversaries.

India's threat landscape is specific and well-documented. State-sponsored groups operating from China and Pakistan have demonstrated sustained interest in Indian critical infrastructure — power grids, banking systems, government networks, and telecom. Ransomware operators have increasingly targeted Indian financial institutions and IT service providers, which collectively manage infrastructure for much of the region. Supply-chain attackers targeting Indian IT services firms gain downstream access to the clients those firms manage.

Against these adversaries, the LSE researchers' warning carries particular weight: the Mythos system card confirms a plausible risk that the model could carry out actions leading to substantially higher odds of a global catastrophe [5]. More immediately, it is the first model to complete a 32-step corporate network attack simulation autonomously — the kind of multistep operation that previously required a coordinated team. When equivalent capability reaches adversarial hands — through state development, through leaked weights, through a less safety-

Restricted release is a first step, not a strategy. The question is whether states cooperate before capability spreads or scramble to respond after.

— Dr Beatriz Lopes Buarque & Abdullah Abu-Hassan, Media@LSE, May 2026 [5]

conscious competitor releasing an open model — the attack surface facing Indian enterprises will have changed in kind, not just degree.

Anthropic is candid on this timeline: models as capable as Mythos Preview will soon be developed by many different AI companies [1]. The question for Indian CISOs is not whether that happens. It is whether Indian enterprises will have hardened their environments in time.

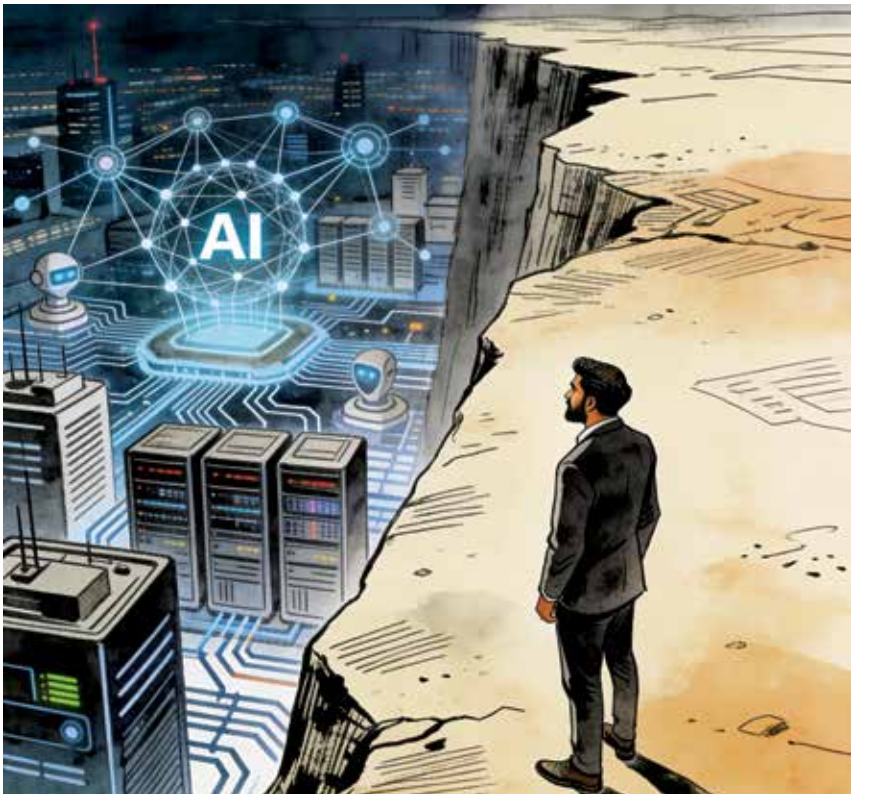

Implications for the CISO Role: A Dimension-by-Dimension Assessment

Threat Intelligence and Risk Assessment

Your threat intelligence baseline has changed. Vulnerabilities that previously required months of skilled manual exploitation research can now be surfaced and weaponised in hours. CERT-In's existing patch windows were calibrated for a world where exploit development was expensive and slow. That world no longer exists.

Mythos Preview demonstrated this concretely: it constructed a working exploit for a vulnerability that had existed in the OpenBSD operating system's codebase — one of the world's most security-conscious projects — for 27 years, undetected [5]. The vulnerability class is not exotic. OpenBSD underlies infrastructure across the Indian banking and government stack. When similar findings from Anthropic's open-source scanning enter the 90-day coordinated disclosure window, they will become public. The patch race begins at that point, and it is a race between defenders who know the window and adversaries who will know it simultaneously.

The immediate action: review every patch SLA for critical and high-severity vulnerabilities. The window that was acceptable under the old threat model is not acceptable now. Any lag must be treated as a residual risk requiring explicit board sign-off and documented compensating controls.

Security Operations: Detection and Response

Adversaries with Mythos-class capability will not attack at human speed. They will deploy autonomous agents that simultaneously operate across multiple hosts and networks without human command [5]. The LSE researchers confirm that Mythos Preview is the first model to complete a 32-step corporate network attack simulation end-to-end — operations that would take a skilled human twenty hours, executed autonomously [5].

SOC workflows calibrated for human-speed attack progressions cannot detect or respond to this threat profile in time. An autonomous attacker can progress from initial access to lateral movement to data exfiltration before a human analyst has finished reading the first alert.

Indian enterprise CISOs — whether running in-house SOCs, hybrid models, or outsourced managed security services — must assess whether their detection logic and response playbooks are designed for this reality.

The compensating controls are not new. They become more urgent. Harden default network configurations. Enforce multi-factor authentication across every privileged access pathway. Maintain comprehensive logs. Implement zero-trust segmentation. These controls buy time when a patch has not yet landed. In a Mythos-class threat environment, time is the only thing that matters [1].

The Defender Dividend: What AI Does for the Blue Team

The picture is not only threat. The same capability that expands adversarial attack surface also gives defenders something genuinely unprecedented: the ability to find vulnerabilities in their own systems

before adversaries do, at a scale and speed no human team could match.

The Glasswing data is compelling on its own terms. Partners report bug-finding rates increasing by more than a factor of ten [1]. At one Glasswing partner bank, Mythos Preview detected and stopped a $1.5 million fraudulent wire transfer after a threat actor compromised a customer email account [1]. Vendors embedded in the Glasswing and OpenAI TAC ecosystems — IBM, Qualys, CrowdStrike, Palo Alto Networks, SentinelOne — are translating these capabilities into enterprise security products that Indian organisations already buy [2][3][6].

The key analytical point from Qualys CTO Dilip Bachwani is worth retaining: the most significant impact of frontier AI may not be the threats it creates, but the speed at which it changes the economics of cyber risk [3]. Traditional vulnerability management was designed for human-speed discovery and human-speed remediation. Both sides of that equation have now changed. Defenders who access AI-assisted scanning gain an asymmetric advantage — but only if they have the human capacity to act on what the AI surfaces.

Vendor and Supply Chain Risk

Frontier AI capability is not only arriving through direct Glasswing or TAC programme access. It is arriving embedded in the security products your vendors are already shipping. IBM Concert, Qualys Enterprise TruRisk Management, CrowdStrike, Palo Alto Networks, SentinelOne, Snyk — these platforms are integrating Mythos-class and GPT-5.5 capabilities as standard features [2][3][6].

This creates a dual obligation. First, understand what frontier AI capabilities are now embedded in your existing security stack, what data those capabilities process, where processing occurs, and what your vendor's human-oversight model looks like. SEBI has explicitly flagged that AI-based vulnerability tools raise concerns around data confidentiality and output reliability [4] — that concern applies to capability embedded in your vendors' platforms, not just tools you procure directly. Second, and more structurally: if your managed security service provider is deploying frontier AI scanning on your environment without a documented validation and escalation process, you are not gaining a defender advantage. You are generating findings that may not be acted on correctly.

Regulatory Compliance and Data Governance

The Indian regulatory environment is moving at pace. SEBI's cyber-suraksha.ai task force has issued guidance covering immediate patching requirements, AI-based vulnerability assessment processes, stronger vendor coordination, stricter API controls, enhanced monitoring, zero-trust frameworks, and mandatory onboarding onto market security platforms [4]. The RBI is coordinating with global regulators on Mythos-related risks in real time [4].

The Digital Personal Data Protection Act 2023 creates binding constraints on processing personal data with foreign-hosted AI. Any frontier AI security tool that ingests logs, network telemetry, source code, or configuration data from production systems containing personal data triggers DPDPA obligations. The government's sovereign hosting ambition — still unresolved — signals that those constraints may tighten further for regulated sector applications [4].

CISOs in banking, insurance, capital markets, telecom, and healthcare should treat regulatory mapping as a prerequisite to procurement, not a post-implementation compliance exercise. The mapping must cover DPDPA 2023, sector-specific RBI and SEBI cybersecurity directions, and applicable MeitY guidelines. Engage your legal and compliance teams now. The regulatory framework is being written around decisions you are making this quarter.

Board Reporting and Cyber Risk Governance

Frontier AI creates a reporting obligation and a reporting opportunity simultaneously. The obligation: your board needs to understand that the threat environment has materially changed, that existing patch timelines and SOC playbooks may be calibrated against a threat model that is already obsolete, and that the regulatory landscape is shifting in real time. Finance Minister Sitharaman and IT Minister Vaishnaw flagged the Mythos risk publicly in April [7]. If the finance minister has flagged it, your board should already have been briefed.

The opportunity: CISOs who can frame the defensive case clearly — what AI-assisted scanning could surface in your environment, what closing those vulnerabilities is worth in reduced breach risk, what the cost of not acting is — will strengthen their position as strategic advisors rather than operational managers.

Frame the board conversation around three questions. What is our current exposure to vulnerabilities that a Mythos-class adversary could find and exploit? What is our plan and timeline to close the most critical gaps? And what is our path to accessing defensive AI capability proportionate to the threat we now face?

Talent, Skills, and the AI Security Analyst

The skills required to deploy, validate, and govern frontier AI security outputs do not map onto existing job descriptions. The analyst this era requires is part vulnerability researcher, part model-output critic, part governance specialist. They must be able to recognise a plausible hallucination in a triage

finding, contextualise AI output against a live organisational threat model, and document the human-in-the-loop decision trail that regulators will eventually require.

This profile is scarce everywhere. It is acutely scarce in India. The immediate risk for Indian enterprises dependent on managed security service providers is that those providers deploy frontier AI tools at volume without having built the internal expertise to validate what those tools produce — generating noise that looks like signal, and signal that gets lost in the noise.

The near-term investment priority is targeted upskilling: vulnerability classification, AI output validation, prompt engineering for security workflows, and human-in-the-loop governance design. This is not a training-course purchase. It is a capability-building programme, and it needs to start before the tools are deployed, not after.

The Cost and Skills Reality: What Deployment Actually Requires

Access to frontier AI cybersecurity capability is not a procurement decision. It is a systems integration challenge with cost and skills barriers that most Indian enterprises are not currently positioned to clear without deliberate and sustained investment.

The Cost of Running Frontier AI at Security Scale

Frontier AI models operate on usage-credit economics, not annual licences. Anthropic provided $100 million in usage credits to its initial Glasswing cohort of approximately 50 organisations — roughly $2 million per partner for exploratory usage [5]. Production-grade agentic scanning of large codebases at continuous cadence costs significantly more. Token costs for multi-step security reasoning tasks are materially higher than for standard generative AI workloads, and the agentic scanning pipelines that produce the most valuable findings are the most compute-intensive.

The remediation cost compounds this further. Anthropic's own data shows each high- or critical-severity vulnerability found by Mythos Preview takes an average of two weeks to patch [1]. AI finds faster than humans can fix. Enterprises that deploy frontier AI scanning without proportionate investment in remediation engineering capacity will generate finding backlogs they cannot clear — widening the disclosure gap rather than closing it.

The Hallucination Risk: Understanding the Accuracy Picture

Frontier AI models are not infallible. In security contexts, an error is not an inconvenience — it is an operational risk.

Anthropic's open-source scanning data shows that 9.4% of vulnerabilities initially assessed by Mythos Preview as high- or critical-severity were false positives on independent review [1]. That figure deserves both sides of its framing. A 90.6% true-positive rate at high-severity tier is, by any historical benchmark for automated tooling, extraordinary — Cloudflare's team rated it better than human testers [1]. But in a regime where the model surfaces thousands of high-severity findings, a 9.4% false-positive rate still represents hundreds of findings that consume analyst time and, if acted on incorrectly, divert remediation resources from real vulnerabilities. SEBI's explicit concern about output reliability [4] is operationally grounded.

The deployment principle is non-negotiable: AI amplifies analyst capacity; it does not replace analyst judgment. Every finding requires human validation before any remediation action is taken. The skills to perform that validation correctly are the scarcest input in the entire pipeline.

Key Deployment Risks for Indian Enterprises

- **Access lag:** Indian organisations entered Glasswing in the second wave. Six to eight weeks of vulnerability scanning data and patching experience that founding-cohort organisations have accumulated does not yet exist for Indian enterprises [7].
- **Sovereign hosting delay:** The government's ambition for domestic hosting of Claude AI creates a compliance constraint for regulated sectors that limits production-scale deployment until the framework is resolved [4].
- **Cost structure:** Agentic scanning at enterprise scale represents a fundamentally different cost profile from annual-licence security tooling. Remediation capacity must scale with discovery capacity or backlogs widen [1].
- **Hallucination in triage:** A 9.4% false-positive rate at high-severity tier (Anthropic's own data) requires skilled human validation at every step — a capability in short supply [1].
- **Remediation bottleneck:** Average two weeks per high-severity patch [1]. AI accelerates discovery far faster than human teams can fix, creating a disclosure gap adversaries can exploit.
- **Skills deficit:** The AI security analyst profile required to govern frontier AI security tools does not exist at scale in the Indian talent market.
- **Vendor opacity:** Frontier AI capabilities embedded in third-party security platforms may not be transparently disclosed, governed, or validated — creating invisible hallucination risk in your existing stack.

Strategic Recommendations for Indian Enterprise CISOs

The following recommendations are sequenced by urgency. Some require board sponsorship; all require action.

Immediate (0–90 Days)

- Review every patch management SLA for critical and high-severity vulnerabilities. The window acceptable under the old threat model is not acceptable now. Treat any gap as a residual risk requiring board sign-off and documented compensating controls.
- Map your open-source and third-party dependency exposure against the Glasswing vulnerability class. Operating systems, web browsers, cryptographic libraries, and network infrastructure libraries are all in scope. Findings will enter the public domain as the 90-day coordinated disclosure window closes — you need to know your exposure before they do.
- Interrogate your top five security vendors on frontier AI integration. Which capabilities have they embedded? What data do those capabilities process? Where is processing performed? What is their hallucination mitigation and human-override model? If they cannot answer clearly, that is itself a risk finding.
- Complete a DPDPA and sector-regulatory mapping before any AI security procurement. For banking, insurance, capital markets, telecom, and healthcare CISOs, this mapping is a compliance prerequisite. Engage legal and compliance teams immediately. The regulatory framework is being written around decisions you are making this quarter.
- Brief your board now, if you have not already. Finance Minister Sitharaman flagged this risk publicly in April [7]. Your board should not be hearing about frontier AI cybersecurity risk from a newspaper.

Medium Term (90–180 Days)

- Pilot AI-assisted vulnerability scanning in a bounded, non-production environment. Use generally available frontier AI models — Claude Opus 4.7 or GPT-5.5 with Trusted Access for Cyber — with explicit human validation at every triage step. Document false-positive rates and remediation timelines before scaling. Establish baseline metrics before the programme expands.
- Invest in targeted analyst upskilling: vulnerability classification, AI output validation, prompt engineering for security workflows, and human-in-the-loop governance design. This capability must precede deployment, not follow it.
- Review SOC detection logic and incident response playbooks against machine-speed attack progressions. Identify the specific gaps where human-in-the-loop response timelines will fail against an autonomous adversary. Prioritise automated detection in those gaps.
- Engage CERT-In, SEBI, and RBI proactively on your organisation's posture and approach to frontier AI security tools. Early engagement provides regulatory visibility and positions your organisation as a responsible actor in a framework that is still being designed.

Strategic (180 Days and Beyond)

- Develop a frontier AI security governance framework: procurement criteria, deployment standards, validation requirements, human-override protocols, audit trails, and explainability obligations. Design it for the regulatory environment you expect in 18 months, not the one that exists today.
- Engage actively with cyber-suraksha.ai and CERT-In

CISO Self-Assessment: Frontier AI Readiness

- Have you reviewed patch management SLAs against machine-speed exploit development timelines?
- Do you have a current inventory of open-source and third-party dependencies mapped against Glasswing-class vulnerability categories?
- Have you interrogated your top five security vendors on what frontier AI capabilities are now embedded in their platforms and how data is governed?
- Have you completed a DPDPA and sector-regulatory mapping for any planned AI security tool data flows?
- Does your SOC have detection logic and playbooks designed for autonomous, machine-speed attack progressions?
- Do your security analysts have the skills to validate and contextualise frontier AI security outputs, including recognising likely hallucinations?
- Have you briefed your board on the changed threat environment and your organisation's readiness posture?
- Are you engaged with CERT-In, SEBI, or RBI on your frontier AI security approach?
- Do you have a documented plan for accessing defensive AI capability proportionate to the adversarial threat?

Score: 8–9 yes — strong readiness posture. 5–7 — material gaps; structured remediation required. 4 or below — immediate board escalation advised.

coordination initiatives. India's regulatory and incident-response infrastructure for frontier AI security is being architected now. CISOs who participate help shape it. Those who wait inherit whatever was built in their absence.

- Build for capability parity, not vendor dependence. The strategic goal is an organisational capability to evaluate, deploy, validate, and govern whatever frontier AI security tools are available — not a dependency on a single vendor whose access terms, pricing, or regulatory status may change. Vendor lock-in in this space carries strategic risk.

Conclusion: The Interim Period Is the Danger Zone

The arrival of Claude Mythos and GPT-5.5 Cyber does not create a future problem. It creates a present one. The specific asymmetry at its core: the ability to find vulnerabilities has accelerated by an order of magnitude; the ability to patch them has not. Anthropic acknowledges the bottleneck directly — it is human capacity to triage, report, and deploy fixes [1]. That constraint is identical for Indian enterprises.

India's position adds a second asymmetry. The founding Glasswing cohort spent April and May closing vulnerabilities in software that Indian enterprises depend on. India's government agencies, armed now with access to Mythos Preview, will begin the same work — but weeks behind, under sovereign hosting constraints that complicate regulated-sector deployment, and without the established remediation pipelines that the founding cohort has had time to build.

The interim period — while vulnerabilities are being rapidly discovered and slowly patched, while defenders are building capability and adversaries are working to acquire equivalent access, while regulatory frameworks are still being written — is the period of highest risk. The decisions Indian enterprise CISOs make in the next ninety days will shape their organisations' exposure posture for years.

The tools themselves are extraordinary. But tools do not govern themselves. The human judgment required to deploy them responsibly, validate their outputs rigorously, navigate the regulatory landscape taking shape around them, and translate findings into board-level decisions that get funded — that is the CISO's work. In the frontier AI era, it has never been more consequential. ■

References and Sources

[1] Anthropic. "Project Glasswing: An Initial Update." Anthropic Research. May 22, 2026. https://www.anthropic.com/research/glasswing-initial-update

[2] IBM Newsroom. "IBM Brings Its Most Advanced AI-Powered Security Portfolio to Clients, and is Strengthened by Ongoing Project Glasswing Work." May 19, 2026. https://newsroom.ibm.com/2026-05-19-IBM-Brings-Its-Most-Advanced-AI-Powered-Security-Portfolio-to-Clients,-and-is-Strengthened-by-Ongoing-Project-Glasswing-Work

[3] Bachwani, Dilip. "Advancing Cybersecurity in the Age of Frontier AI: Qualys Steps into Project Glasswing." Qualys Blog. June 5, 2026. https://blog.qualys.com/product-tech/2026/06/05/advancing-cybersecurity-in-the-age-of-frontier-ai-qualys-steps-into-project-glasswing

[4] Bisht, Shrishti. "India Gets Access to Anthropic's Mythos AI Model Under Project Glasswing." Inc42. June 3, 2026. https://inc42.com/buzz/india-gets-access-to-anthropics-mythos-ai-model-under-project-glasswing/

[5] Lopes Buarque, Beatriz and Abu-Hassan, Abdullah. "Claude Mythos and the Myth of Containment." Media@LSE, London School of Economics. May 11, 2026. https://blogs.lse.ac.uk/medialse/2026/05/11/claude-mythos-and-the-myth-of-containment/

[6] OpenAI. "Scaling Trusted Access for Cyber with GPT-5.5 and GPT-5.5-Cyber." OpenAI News. May 7, 2026. https://openai.com/index/gpt-5-5-with-trusted-access-for-cyber/

[7] Barik, Soumyarendra. "Project Glasswing: Key Cybersecurity Agencies Set to Get Access to Anthropic's Mythos." The Indian Express. June 6, 2026. https://indianexpress.com/article/business/companies/project-glasswing-key-cybersecurity-agencies-set-to-get-access-to-anthropics-mythos-10727692/

Cybercriminals already scored before the whistle blows

FortiGuard Labs maps a sophisticated cybercrime operation — fake domains, stolen credentials, and malware — already targeting FIFA 2026.

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

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

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 week 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,

Before the first whistle blows at FIFA World Cup 2026, cybercriminals have already kicked off their own tournament.

then directed them to credential-harvesting pages that mimicked Google login. A shared Google Analytics ID across dozens of these domains' points to a single organized threat actor running the operation at scale.

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 about 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. ■

Asia-Pacific's cybercrime boom: Inside INTERPOL's 2025/2026 threat report

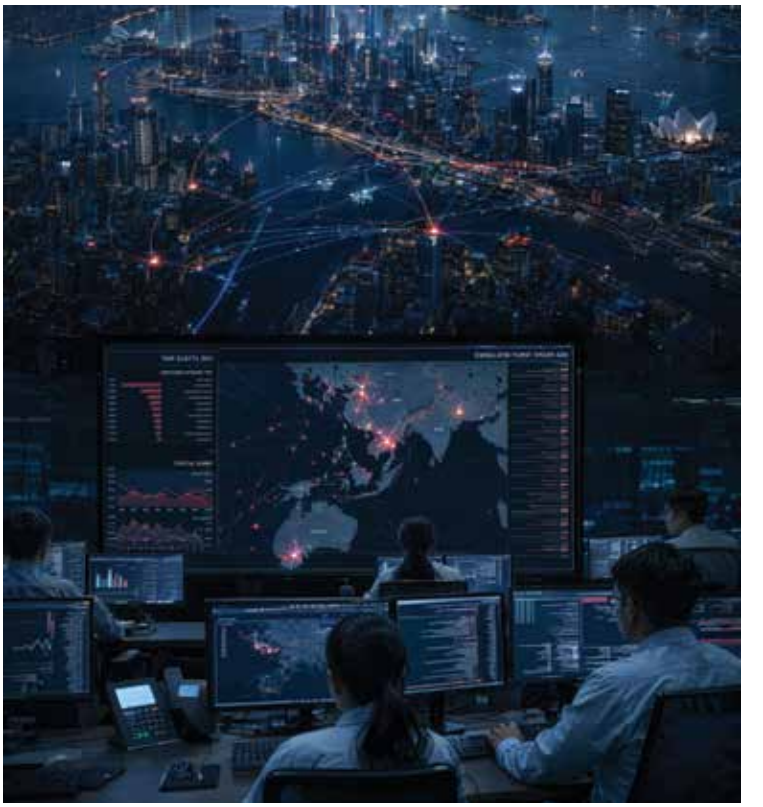

INTERPOL's 2025/2026 report reveals a surge in AI-driven scams, ransomware, and infostealer malware across Asia-Pacific's digital economy.

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

A new report from INTERPOL paints a stark picture of the Asia and South Pacific region: a digital economy outpacing its defenses. The INTERPOL Asia and South Pacific Cyber Threat Assessment Report 2025/2026, prepared by the Asia and South Pacific Desk (ASP Desk) under the ASPJOC initiative, draws on survey responses from 18 member countries along with private-sector intelligence to map out how cybercrime is evolving across the region between January 2024 and March 2025.

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.

Scams are the biggest problem

Online scams and phishing top the list of threats, with a third of surveyed countries reporting over 10,000 cases each. These aren't small-time operations — organized crime groups in Cambodia, Lao PDR, Myanmar, and the Philippines have built industrial-scale scam centers, some resembling forced labor camps, generating an estimated USD 40 billion annually through romance scams, fake investments, and illegal gambling.

Ransomware and malware remain relentless

The region logged more than 135,000 ransomware attacks in 2024 alone, hitting real estate, manufacturing, and finance hardest. One attack on Indonesia's National Data Center knocked out 280 essential government services. Meanwhile, infostealer malware — tools like RedLine, LummaC2, and Lokibot — quietly harvest passwords and financial data, often setting the stage for bigger fraud or ransomware attacks down the line.

AI is changing the game for criminals too

Perhaps the report's most alarming finding is how fast criminals are adopting AI. Deepfake discussions on criminal forums jumped 600% in just a few months of 2024. Real-world consequences followed: a Hong Kong firm lost USD 25 million after a deepfake video call impersonated executives, and a Singapore finance director nearly lost half a million dollars the same way. AI is also powering more convincing phishing emails, fake job scams, and malware that can slip past traditional security tools.

Cybercrime in this region is no longer scattered or amateur — it's organized, technologically sophisticated, and growing.

The numbers are staggering

Over 6.5 billion cyber threats were blocked in the region during 2024, most of them arriving via email. Data breaches remain common, with system intrusions accounting for 80% of cases.

Fighting back

The good news: countries aren't standing still. February 2025's Operation SECURE united 26 nations to dismantle infostealer infrastructure, seizing servers and taking down over 20,000 malicious domains. Nations like South Korea, Indonesia, Hong Kong, and Japan have rolled out new laws, AI-powered detection tools, and public awareness campaigns. Notably, two-thirds of the surveyed law enforcement agencies have already adopted AI.

The takeaway

INTERPOL's message is clear: cybercrime in this region is no longer scattered or amateur — it's organized, technologically sophisticated, and growing. Closing the gap will require stronger cross-border cooperation, better-resourced law enforcement, and closer partnerships with the private sector. ■

AI's new battleground: Infrastructure, not algorithms

AI sovereignty is unrealistic; alone, nations must balance domestic control with trusted global partnerships.

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

A new World Economic Forum white paper, produced by Bain & Company, argues that AI infrastructure, not algorithms, has become the real battleground for national competitiveness. Titled "AI Infrastructure in the Age of Sovereignty," the report lays out how nations should think about building, sharing, or borrowing the physical backbone that powers AI.

The three pillars of AI infrastructure

The report breaks AI infrastructure into three building blocks: compute (processing power), connectivity (networks linking data), and data storage. Each face hard physical limits of energy, water, land, hardware, and cybersecurity that determine what's buildable. Global data center electricity use alone could nearly triple by 2035, and a single AI training facility can require 200 acres of land and millions of gallons of water daily.

No country can go it alone

Perhaps the report's central message is that true self-sufficiency in AI infrastructure is unrealistic for almost everyone. Even China and the US, the only economies with near-full domestic control, still rely on concentrated global supply chains for chips. Most nations, the report says, must choose a strategy somewhere between two extremes:

- **Trusted international partnerships —** leaning on foreign partners for compute and storage while keeping legal control (think Estonia storing government data in Luxembourg).
- **Extensive domestic ownership —** building and controlling everything at home, best suited for well-resourced economies or sensitive workloads like defense.

In practice, most countries will blend both, as Singapore does by combining a domestic supercomputing center with heavy use of international cloud providers.

Five things need to be trusted: political commitment, legal clarity, data governance rules, technical safeguards, and transparent operations.

Enter "Digital Embassies"

The report's most novel contribution is a Global Framework for Innovative and Trusted Digital Embassies. Secure, extraterritorial digital spaces in which a country's data and AI workloads are hosted abroad yet remain governed by its own laws. Pioneered by Estonia's 2017 data-hosting deal with Luxembourg, this model is now gaining traction, with Saudi Arabia and Bahrain drafting similar legal structures.

For this work, the Forum argues that five things need to be trusted: political commitment, legal clarity, data governance rules, technical safeguards, and transparent operations. Without them, digital embassies risk becoming legal grey zones rather than genuine extensions of sovereignty.

The bottom line

The report's call to action is blunt: governments, investors, and infrastructure operators must treat energy and land constraints as design realities, not afterthoughts, and build partnerships with enforceable safeguards rather than blind trust. As AI infrastructure decisions today will shape economic resilience for decades, the paper positions digital embassies and disciplined, hybrid strategy-making generally as key tools for countries seeking sovereignty without isolation. ■

Ransomware in 2026: Costlier attacks, smarter defences

Sophos's 2026 report shows ransomware attacks growing costlier and more frequent, even as victims pay less and recover faster.

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

A new report from Sophos, The State of Ransomware 2026, paints a picture of a threat landscape in flux. Based on responses from 2,158 IT and security leaders across 17 countries, the seventh annual survey shows organisations paying less to attackers, recovering data faster, and negotiating harder even as the overall cost and frequency of successful breaches continue to climb.

Email, not bugs, is the new front door

For three years running, exploited software vulnerabilities were the top cause of ransomware attacks. That's no longer true. This year, malicious email (26%) and phishing (24%) overtook technical exploits (18%) as the leading entry points. Compromised credentials remained a steady factor at 23%. Altogether, 79% of attacks began with some form of identity-based intrusion, stolen or abused login credentials making identity security, not just patching, the new frontline of defence.

Reinforcing this, two-thirds of victims (67%) said their ransomware incident was the very same event as their most significant identity breach. Even more strikingly, 97% of organisations hit through compromised credentials already had multi-factor authentication turned on somewhere in their environment, proof that MFA alone isn't a silver bullet unless it's deployed comprehensively.

Attacks are succeeding more often

After two years of improvement, encryption success rates ticked back up to 56%, from a low of 50% in 2025. Sixteen per cent of victims had data both encrypted and stolen from a dangerous combination that gives criminals extra leverage. Firewalls proved valuable early-warning tools: when they flagged an attack before ransomware was deployed, encryption succeeded only half the time, compared with 71% when firewalls missed it entirely.

Smaller payments, faster recoveries

The financial picture is more encouraging. The median ransom demand fell sharply to $698,000, down 65% over two years, and the median payment dropped to $769,000. Victims are also getting tougher negotiators; 51% paid less than the original demand. Backup-based recovery surged to 66% of encrypted data cases, up from 54% last year, while just 48% of victims felt compelled to pay a ransom at all, the lowest rate in three years.

Yet recovery itself isn't cheap: the average cost to clean up after an attack, excluding any ransom, hit $1.7 million, an 11% jump from 2025.

The human toll

Perhaps the report's most sobering finding: 99% of organisations that suffered an encryption incident reported lasting effects on their security teams, ranging from rising anxiety (41%) to leadership shake-ups (21%). One bright spot: recognition from senior leadership rose to 36%, suggesting boardrooms are finally paying closer attention.

Sophos's takeaway is clear: as AI accelerates both attacks and defences, organisations that unify identity protection, email security, and backup resilience rather than treating them as separate problems will be best positioned for the fights ahead.■

Are you aligned with your CDO for AI strategy?

If AI is becoming everyone's business, data can no longer be one function's responsibility. And if data is becoming fuel for AI, the CISO and CDO may need to be more closely aligned than ever before.

I have been thinking about this increasingly through my conversations with CISOs. One of the biggest challenges organizations face today is building a data strategy that allows AI to scale without compromising security, privacy, or trust.

As AI moves deeper into enterprise operations, and AI agents begin retrieving information, making decisions and taking actions, the CISO faces a more nuanced question: Which data is truly critical? Which is sensitive? Who should have access to it, and under what circumstances can an AI system use it?

This is where the CISO-CDO relationship becomes critical.

The CDO sees data as a strategic asset, to be connected, enriched, and leveraged for AI, analytics, and business outcomes. The CISO looks at the same data through the lens of risk: where it resides, who can access it, how it can be exposed, and what happens if it is misused.

Consider a bank deploying an AI agent for relationship managers. The CDO may want it to access customer profiles, transactions, and interaction histories to provide contextual responses. The CISO must ask: Should the agent have access to all of that information? Can it retrieve data beyond the employee's permissions? Can sensitive information be exposed through an AI response?

The same challenge exists in manufacturing, where AI may need operational and supplier data but could also expose commercially sensitive information or vulnerabilities in critical infrastructure.

Neither perspective can work in isolation.

As Marc Hogan, Contributor, IBM Think Leadership, has noted, "The cost of a fractured CDO-CISO relationship? Blocked innovation on one side, and compliance risks on the other."

IBM's 2025 Chief Data Officer Study reinforces the gap: 87% of executives say effective data security is vital to realizing AI investments, but only half say they can protect sensitive data in most AI use cases.

This is a conversation we need to have more often. Are CISOs and CDOs jointly deciding what constitutes critical data? Are security and privacy controls being built into AI architecture from the beginning?

As AI moves from copilots to agents, the CISO-CDO relationship could determine whether enterprises can move fast and remain secure. ■

"One of the biggest challenges organizations face today is building a data strategy that allows AI to scale without compromising security, privacy, or trust."

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

AI IS SCALING.

CAN YOUR INFRASTRUCTURE KEEP UP?

ET Edge | CIO&LEADER

AI INFRASTRUCTURE SUMMIT 2026

THE INFRASTRUCTURE BEHIND THE AI ENTERPRISE

16 OCTOBER 2026 | MUMBAI

COMPUTE | CLOUD | DATA | NETWORK | POWER | SECURITY

India's top CIOs come to decode what it takes to scale AI.

From the compute that powers AI to the cloud, data, network, power and security that sustain it, every layer of the enterprise stack is being redefined.

75+ INDIA's TOP CIOs	12+ INSIGHT-LED SESSIONS	EXCLUSIVE CIO WORKSHOPS	CLOSED-DOOR ROUNDTABLES	HIGH-VALUE CXO NETWORKING

BUILD WHAT AI NEEDS TO SCALE.

THIS IS NOT A CONFERENCE ABOUT TALKING ABOUT AI.
THIS IS ABOUT DECIDING WHAT NEEDS TO SIT UNDERNEATH IT.

For further information, kindly write to

Hafeez Shaikh
hafeez.shaikh@timesgroup.com
+91 9833103611

Aanchal Gupta
aanchal.gupta1@timesgroup.com
+91 9651841119

Supriya Sahoo
supriya.sahoo@timesgroup.com
+91 8095056886

Sidharth
sidharth@timesgroup.com
+91 9910386143

What if 45 minutes could change your ciso relationships?

ET Edge CISOFORUM

CISO PLAYBOOK 2026

Reset. Architect. Defend.

16 OCTOBER 2026 **MUMBAI**

AI Security | Cloud & Identity | Data Protection | Cyber Resilience

The next generation of enterprise security won't be shaped by presentations. It will be shaped by the right conversations.

One partner. One think tank.
Move beyond visibility. Create influence.

Ciso playbook is built for the brands shaping the next security era

Own the CISO conversation.

For further information, kindly write to

Hafeez Shaikh
hafeez.shaikh@timesgroup.com
+91 9833103611

Aanchal Gupta
aanchal.gupta1@timesgroup.com
+91 9651841119

Supriya Sahoo
supriya.sahoo@timesgroup.com
+91 8095056886

Sidharth
sidharth@timesgroup.com
+91 9910386143