

CISOFORUM

Security For Growth And Governance

THE IDENTITY RECKONING

What Every CISO Must Get Right in 2026



From Left to Right : Vasanth Pai Global DPO, Tech Mahindra | Pradnya Manwar CISO & DPO, Thomas Cook | Baidyanath Kumar CISO & DPO, JK Lakshmi Cement | Sunil Nishankar CISO & DPO, Ageas Federal Life Insurance | Komal Vora CISO, Equifax Credit Information



CISOFORUM



CISO
PLAYBOOK 2026

Reset. Architect. Defend.



27-28 NOVEMBER 2026



KARJAT

AI IS REWRITING THE ENTERPRISE.

WHO'S REWRITING THE SECURITY PLAYBOOK?

75+

SENIOR CISOS



AI Security



Cloud & Identity



Data Protection



Cyber Resilience

TWO
IMMERSIVE
DAYS

LEAD THE THINKING | DEMONSTRATE WHAT'S NEXT | DEEPEN THE CONVERSATION.

BEYOND THE NOISE. INTO THE CONVERSATION.

GET INTO THE PLAYBOOK.
WRITE TO

Hafeez Shaikh

hafeez.shaikh@timesgroup.com
+91 9833103611

AN EVENT BY



THE TIMES GROUP

INDIA'S LEADING B2B INTELLIGENCE, COMMUNITY & MARKET PLATFORM

Is Identity Running Out of Time?

Identity and access management (IAM) was built on three assumptions: that identities are human, that attackers need time, and that the mathematics beneath credentials is permanent. Frontier AI and quantum computing are eroding all three. If you treat them as separate agendas, then you will miss what they share. One attacks the implementation of trust; the other attacks its mathematics.

In May, Anthropic reported that it had used Claude Mythos Preview to find more than ten thousand high- or critical-severity vulnerabilities within weeks of launch. One flaw found in wolfSSL, a widely embedded cryptography library, would have let attackers forge certificates. Cloudflare, a partner, observed that the model chains low-severity bugs, once left in backlogs, into severe exploits. Anthropic cautioned that comparable models will soon be widely available and urged defenders to favor controls that work without waiting for a patch: hardened defaults, multi-factor authentication, comprehensive logging. When exploitation is cheap, identity becomes the containment layer.

NIST's National Cybersecurity Center of Excellence has proposed adapting existing identity standards, rather than inventing new ones, for AI agents. A joiner-mover-leaver process was never designed for an identity created in seconds, acting at machine speed and delegating to others. Agents need owners, narrow scopes, and short-lived credentials from birth.

Quantum is the second clock. When Google set a 2029 migration deadline in March, it adjusted its threat model to prioritize authentication services. Encryption is already exposed to harvest-now, decrypt-later collection, but signatures must be replaced before a capable machine exists. Identity runs on signatures: certificates, token signing, code signing and passkeys, which today rely on elliptic-curve or RSA keys. NIST's proposes deprecating by 2030. India's task force roadmap sets a December 2029 for critical infrastructure, banking included.

For CISOs, three actions are imperative. Inventory every point where your identity fabric depends on RSA or elliptic curves, and demand crypto-agility from the passwordless programs you are using today. Treat every agent as an identity: registered, owned, scoped, expiring. And assume vulnerabilities beneath your directory will be found faster than patched, so a stolen credential should have low impact.

Identity has always asked who, or what, may be trusted. Two recent developments are now shortening the answer's shelf life. ■



“Identity has always asked who, or what, may be trusted. Two recent developments are now shortening the answer's shelf life.”

R. Giridhar

Editorial Director

Enterprise Tech Publications

c-raja.giridhar@timesgroup.com

CONTENTS

SEPTEMBER 2026



COVER STORY

09-19

The Identity Reckoning

Three regulators are enforcing four regimes that converge on the same demand, even as machine identities proliferate and 2026 breach data shows stolen credentials turn up in almost every kind of attack. Yet the control behind all three is rarely anyone's whole job.



Cover Design by:
Manish Kumar



Please Recycle This Magazine And
Remove Inserts Before Recycling

COPYRIGHT All rights reserved:

Reproduction in whole or in part without written permission from 9.9 Group Pvt Ltd (formerly known as Nine Dot Nine Mediaworx Pvt Ltd). Published at 121, Patparganj, Mayur Vihar Phase-1, Near Mandir Masjid, Delhi-110091 and printed at G. H. Prints Private Limited, A-256 Okhla Industrial Area, Phase-I, New Delhi - 110020.

NEWS & VIEWS



05

F5 speeds up cyber defence as AI turbocharges attacks



08

Palo Alto brings OpenAI's cyber models into the defender's toolkit

INTERVIEW



20-26

Machine Speed, Human Stakes: Why India's Cyber Defences Are Running Out of Time
By Jagrati Rakheja



27-29

Data is the new battleground, not the background:
By Jagrati Rakheja

INSIGHTS



30-31

RBI's new cybersecurity rulebook: From advice to law



34-35

Phishing gets smarter, not louder: Inside Zscaler's 2026 threat report

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

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**



Sammeer Kulkarni joined Bajaj Life Insurance as CISO

Sammeer Kulkarni has joined Bajaj Life Insurance as Chief Information Security Officer (CISO). With over 27 years of experience across cybersecurity, information security governance, cyber risk and regulatory compliance, he previously served as CISO at Zurich Kotak General Insurance. He held senior roles at BNP Paribas, IndusInd Bank and IndiaFirst Life Insurance.



Yogesh Kulkarni joined STL as Group CISO – Global

Yogesh Kulkarni has joined STL – Sterlite Technologies as Group Chief Information Security Officer – Global. He previously held CISO roles at Crompton Greaves Consumer Electricals and Pidilite Industries, besides security leadership positions at Lupin, WNS Global Services and Reliance Communications. At STL, he will focus on cybersecurity strategy, resilience, risk management and digital trust.



Vijay Thavasi joins Blackstone as SVP – Cybersecurity Head, APAC

Vijay Thavasi has joined Blackstone as SVP – Cybersecurity Head, APAC. He will lead regional cybersecurity strategy, cyber risk, security operations, cloud security, AI governance, and resilience. With over 17 years of experience, he previously held leadership roles at Rakuten, Williams Lea Tag, Tata Communications and Wells Fargo.



Shrikant Iyer joined Aditya Birla Capital as the Head of Cyber Security

Shrikant Iyer has joined Aditya Birla Capital as Head of Cyber Security. With over 28 years of experience in cybersecurity, information security, governance, risk and compliance, he previously served as CISO at Aditya Birla Health Insurance and Aditya Birla Insurance Brokers. He held senior roles at BSE India and HDFC Bank.

F5 speeds up cyber defence as AI turbocharges attacks

F5 launches AI-powered WAF upgrades and faster virtual patching to help counter exploits that AI can create within hours.

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

Cybersecurity firm F5 has rolled out new tools designed to help companies respond to a growing threat: attacks powered by frontier AI that can turn a software flaw into a working exploit within hours. The company's answer is a mix of smarter detection and much faster patching, built directly into the traffic path where attacks first appear.

The problem: Patching can't keep up

Traditionally, fixing a security flaw meant rewriting code, testing it, and pushing out a new release, a process that can take weeks. F5 argues that timeline no longer works when attackers, aided by AI, can weaponise a vulnerability almost as soon as it's discovered.

"Frontier AI has collapsed the time between vulnerability discovery and active exploitation," said Kunal Anand, Chief

Product Officer at F5. He added that the company's goal is to buy security teams something in short supply: time to assess risk properly instead of scrambling under pressure.

What's new?

F5's AI-powered web application firewall (WAF) now includes anomaly detection that learns each application's normal traffic patterns and flags unusual activity that could signal an attack, rather than relying only on known threat signatures. That approach aims to catch zero-day attacks and exploits that constantly change shape to dodge detection.

A second addition, agentic threat intelligence, builds on F5's acquisition of Fletch. It combines outside intelligence on active threats with what F5 observes on a customer network, giving security teams a clearer picture of which alerts matter and suggested fixes they can apply immediately.

Patching in minutes, not weeks

Perhaps the most practical upgrade is "virtual patching," which lets teams block a known exploit path at the network level within minutes, buying time until developers can ship a permanent fix. F5 says internal testing showed the WAF caught 98% of threats while keeping false alarms to just 1%, a figure that matters because false positives often drive companies to leave firewalls in passive, non-blocking mode.

Why it matters?

As AI shortens the gap between vulnerability discovery and exploitation, F5 is betting that speed and precision, not just stronger defences, will define the next phase of enterprise cybersecurity. The takeaway is clear: teams need faster detection and virtual patching to stay ahead of AI-driven attacks. The new WAF features are available now, with broader rollout continuing over the coming months. ■





Schools under siege: education sector faces surging email attacks

Education institutions face frequent phishing and slow recovery; new Barracuda survey finds.

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

Educational institutions are dealing with phishing emails and struggling to recover once attackers get in, according to a new global survey from cybersecurity firm Barracuda Networks.

Attacks arrive by the thousands

Barracuda's threat research found that a sample of 536 education institutions received 58.8 million phishing emails over three months—roughly 1,200 per institution every day. As attackers lean more on AI to sharpen their phishing tactics, security teams are finding it harder to catch every malicious message before it lands.

The survey, which covered 2,000 IT and security leaders

across the US, Europe and Asia Pacific, included 113 respondents from the education sector. Of these, 77% said their institution had suffered an email security incident in the past year. Only 34% said they had successfully detected and blocked phishing attempts, and education respondents were four times more likely than the global average to say they couldn't even tell whether an incident had occurred, underscoring how hard it is to spot attacks.

Ransomware and account takeovers hit harder

Education tied for the highest rate of email-driven ransomware globally, at 38% compared with 28% overall, and recorded the highest rate of account takeover, at 27% versus 20% globally.

Recovery lags behind

The sector fell behind the global average at every stage of incident response detection, investigation, containment and recovery, with the widest gaps in investigation and recovery, each trailing by 8 percentage points. Nearly one in five institutions took up to a week to restore normal operations, double the industry average, showing how slowly recovery can stall operations.

Skills gap compounds the problem

Over a third of respondents cited a lack of incident-response expertise as a challenge, while 33% said responding quickly during live incidents was difficult. Just over half pointed to user behaviour and human error as a leading contributor to incidents.

Notably, 41% of education respondents said they wanted trustworthy AI-assisted detection and investigation tools—the highest demand of any industry surveyed—signalling that institutions see these tools as the clearest path to relief. ■

Sophos taps OpenAI's cyber models to tell defenders which vulnerabilities matter

Sophos is adding OpenAI-powered exploit verification to help security teams prioritize which vulnerabilities attackers can reach.

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

Sophos has unveiled a new tool called Exploit Path Verification (EPV), designed to solve one of cybersecurity's most persistent headaches: too many vulnerabilities and not enough clarity on which ones put a company at risk. The capability will be added to Sophos Managed Risk and will run on OpenAI's GPT cyber models through the Daybreak Defence Network. No launch date has been set yet.

The problem with severity scores

Vulnerability scanners are good at finding flaws for thousands of them, but a generic severity score doesn't say whether a

hacker can reach that flaw inside a specific company's systems. A "critical" bug might sit behind a firewall that blocks it entirely, while two "low risk" issues could chain together into a real breach path. Security teams end up patching by score rather than by actual exposure, wasting time on threats that were never reachable.

What does EPV do?

EPV analyses a company's patch status, security controls, network setup, user privileges, and known exploit activity, then sorts each vulnerability into one of four categories: Confirmed Exploitable, Blocked by a Control, Not Reachable, or Insufficient Evidence. It also flags when multiple minor issues combine into a serious attack route, checks whether a defence stops an attack technique or just a known demo exploit, and drafts remediation notes teams can drop straight into a ticket.

Humans still check the work

Sophos says every AI-generated verdict comes with visible evidence and is reviewed by its own analysts before reaching customers; the system is designed to assist, not replace, human judgment. "Exploit Path Verification is being built to make it clear what in their environment is reachable by an attacker, with the evidence to prove it, so they fix what counts first," said John Peterson, Sophos's chief technology officer.

Part of a bigger OpenAI partnership

The move builds on Sophos's membership in OpenAI's Daybreak Defence Network, joined in June 2026, which already feeds frontier AI models into its threat investigations. Sophos protects more than 625,000 organisations and 40,000 managed detection and response customers, and says the goal is to make verified exploit analysis available beyond just the largest, best-funded security teams. ■





"Frontier AI is changing what attackers can find and how quickly exposure can become impact."

operate directly inside customer environments, working under Unit 42's direction to discover exposures, test whether they can be exploited, and validate full attack paths. A multi-model harness routes each task to whichever model performs best, balancing coverage, accuracy and cost. At the same time, human experts continue to verify findings and connect them to real-world threat intelligence.

Why known vulnerabilities aren't enough?

Data backs the urgency behind this shift: 36% of exposures identified by Unit 42 map to no known CVE, often arising from multiple smaller gaps that must be discovered and chained together to be understood as a genuine risk. Traditional vulnerability scanning alone no longer captures the full picture of enterprise exposure.

Speaking on the development, Sam Rubin, SVP, Unit 42 Consulting and Threat Intelligence, shared:

"Frontier AI is changing what attackers can find and how quickly exposure can become impact. By putting OpenAI's specialised cyber models to work for defenders, we are shifting the balance of power. Unit 42 can help organisations proactively simulate adversarial attacks, discover hidden vulnerabilities before they become incidents, and accelerate remediation at machine speed."

The company frames this as part of a broader race against time, arguing that the same AI capabilities reshaping offensive security must now be deployed at equal speed on the defensive side. ■

Palo Alto brings OpenAI's cyber models into the defender's toolkit

Palo Alto now gives defenders direct access to OpenAI's GPT-5.6 Daybreak cyber model via Frontier AI Defence, targeting unknown exposures.

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

Palo Alto Networks has announced a significant expansion of its Frontier AI Defence service, giving security teams direct access to OpenAI's specialised cybersecurity model, GPT-5.6 Daybreak, which has not previously been available for commercial use. The move, powered by a partnership between the two companies, marks a shift in how frontier AI is positioned in the cybersecurity of conversation — not just as a looming threat, but as an active tool for defenders.

From insight to action

Until now, frontier AI models have mostly informed the work of Unit 42 experts. With this expansion, those same models can

THE IDENTITY RECKONING

What Every CISO Must Get Right in 2026

Three regulators are enforcing four regimes that converge on the same demand, even as machine identities proliferate and 2026 breach data shows stolen credentials turn up in almost every kind of attack. Yet the control behind all three is rarely anyone's whole job.

By **R. Giridhar & Jagrati Rakheja** | jagrati.rakheja@timesgroup.com



Ask a CISO who owns identity in their enterprise and listen closely to the pause before the answer. The directory belongs to the infrastructure. Provisioning belongs to IT operations. Access approvals belong to business managers. Joiner and leaver records sit with HR. What the CISO owns is the consequence when any of it fails.

That arrangement survived while identity was plumbing, the unglamorous machinery that let the right people into systems. It stops working when the plumbing becomes evidence. Nobody asks a CISO about the directory itself. They ask whether a particular contractor, service account or software agent could reach a particular system last Tuesday, who approved it, and why it was still switched on.

Most security leaders can answer for a few systems, usually the ones they know best. Very few can answer all of them, on the same day, in a form an auditor will accept. That gap is where attackers operate, and where audit findings turn into board agenda items, usually at the worst possible moment.

Identity is the control every other control quietly depends on. The harder question this year is who in your enterprise can still afford to treat it as somebody else's job.

Identity In Numbers

3.8 — average severity Indian CISOs assign to identity-centric breaches; 68 percent rate the impact high or extreme.

31 percent — rank identity-first Zero Trust architecture as their number one investment priority.

Source: CISO PRIORITIES SURVEY 2026.

Identity is no longer the perimeter

At Gartner's IAM Summit in London in March, one statement resonated across the conference floor: attackers don't break in, they log in. It is still largely true. But 2026 is the first year the industry's own numbers complicate the sentence rather than confirm it. The Verizon 2026 Data Breach Investigations Report (DBIR) — which catalogs more than 22,000 confirmed breaches across 145 countries — found that vulnerability exploitation has overtaken credential abuse as the single largest initial-access vector, 31 percent against 13 percent. It is the first time in the report's nineteen-year history that stolen credentials have lost the top spot.

The headline, though, hides a counting change. Verizon began tracking pretexting as a separate initial access vector this year, absorbing part of what was previously counted as credential abuse; measured the old way; credential abuse stands at 16 percent. Add phishing, also at 16 percent, and identity-led initial access comes to 32 percent — still ahead of vulnerability exploitation. Further along the attack chain, credential abuse appears in 39 percent of breaches, and the human element — phished, socially engineered, or simply careless — sits behind 62 percent.

Identity has not stopped being central to how enterprises get breached. It is the layer that

turns every other foothold — an unpatched server, a leaked session cookie, a manipulated help-desk agent — into a full compromise. The view from inside Indian enterprises matches. In CISO Forum's CISO Priorities Survey 2026, respondents gave identity-centric breaches an average severity score of 3.8, with a 68 percent rating on the impact high or extreme. That place's identity behind ransomware (4.1, 79 percent) and level with insider threats and human risk (3.8, 71 percent).

Consider a composite scenario: a mid-sized Indian BFSI enterprise, midway through a routine internal audit, discovers several hundred dormant service accounts — some carrying elevated database privilege — still active eighteen months after a cloud migration nobody thought to close out. Nobody has exploited them yet. But under three of the four regulatory regimes now in force in India, that audit finding is no longer a private embarrassment. It is a reportable control gap.

CISO Takeaway

Stop measuring identity risk by initial access alone: the 13 percent headline from Verizon DBIR hides a 32 percent identity-led share at the front door and a 39 percent presence across the attack chain.

Four regulatory regimes, one evidence pipeline

The regulatory regimes behind that audit finding trace to three Indian regulators — MeitY, whose remit covers both the DPDP Rules and CERT-In, alongside the RBI and SEBI. Independently, and over the past eighteen months, they have converged on the same demand across four regimes: prove who has access to what and prove it on demand. None of them describe it in an identical language.

CISOs are already feeling pressure. In the CISO Priorities Survey 2026,



"That is the next evolution of cyber resilience: reducing the data-access surface so that an identity compromise does not become a data compromise."

Pradnya Manwar
Chief Information Security Officer
and DPO, Thomas Cook

DPDP Act enforcement is the top-rated regulatory factor — well ahead of sector mandates from the RBI and IRDAI. DPDP compliance also drew 39 percent of first-place priority rankings, second only to resilience at 44 percent.

The DPDP Rules 2025, notified by MeitY on 13 November 2025, name access control, activity logging, encryption and breach-detection evidence as explicit "reasonable security safeguards" under Rule 6 — the first time India's privacy law has made identity controls a codified, auditable obligation. Penalties run up to ₹250 crore for a safeguard failure and ₹200 crore for a delayed breach notification. Compliance is phased across an 18-month window, with most substantive obligations landing by 13 May 2027.

MeitY proposed compressing those timelines at a stakeholder meeting on 22 January 2026, and comments closed on 4 February. As of early September 2026, no amending instrument has been gazetted, so 13 May 2027 remains the operative date. The practical answer for CISOs: plan against eighteen months, and if your organization is likely to be designated a Significant Data Fiduciary, keep a twelve-month contingency plan ready.

On ground, organizational readiness is thin. EY India's January 2026 report, India's Digital Privacy Crossroads, which surveyed around 150 organizations, found 71 percent still struggle to interpret the DPDP Act and Rules with confidence; about 80 percent have not updated their governance policies, and more than 83 percent have not begun implementation. EY's researchers point to five recurring blockers: legacy systems, a shortage of in-house expertise, difficulty interpreting the Rules, cross-border transfer complexity, and budget constraints. At least three of them — legacy systems, missing expertise and interpretation — are problems an identity program runs straight into.

The RBI's Authentication Directions 2025, issued on 25 September 2025, requires two distinct authentication factors, at least one of them dynamic, for domestic digital payment transactions other than card-present ones. The requirement became mandatory on 1 April 2026 and has now been in force for six months. The Directions still expressly permit SMS-based OTP as a factor; the test is whether proof of possession is unique to each transaction, which a per-transaction OTP can satisfy. SMS-OTP is therefore not banned, but banks and fintechs relying on it alone should review whether their implementation meets that test. For high-risk flows, the safer path is device-bound and FIDO-based factors. Cross-border card-not-

present transactions come into scope from 1 October 2026, when issuers must have a validation mechanism and a risk-based mechanism in place for them.

SEBI's Cybersecurity and Cyber Resilience Framework, issued on 20 August 2024 and refined through subsequent technical clarifications, mandates Zero Trust principles, least privilege, MFA, and periodic access review for regulated entities. Compliance ran on a phased schedule — 1 January 2025 for entities already under earlier SEBI cyber circulars, 1 April 2025 for entities newly brought into scope — so the runway has closed for most regulated entities. However, an August 2025 circular issuing technical clarifications (not a new deadline) continues to refine implementation detail.

IRDAI's Information and Cyber Security Guidelines, 2026, issued in April, revise the 2023 framework and apply not only to insurers but to the brokers, corporate agents, web aggregators and third-party administrators that connect to their systems. That reach makes identity a supply-chain question: an insurer's access-review cycle now must account for partner staff it does not employ. The framework being revised already covers privileged access and multi-factor authentication for critical systems, and it held the board ultimately responsible for information security. An insurer therefore answers IRDAI, CERT-In and the DPDP Rules for the same underlying access of evidence.

CERT-In's Directions, which sit under MeitY alongside the DPDP regime, run underneath the others: six-hour incident reporting, 180-day in-India log retention, and annual audits by CERT-In-empaneled firms that increasingly demand hard evidence — log samples, configuration exports, patch records — rather than a signed policy attestation.

Multiple regimes and regulators — one underlying evidence pipeline,



"An AI agent for fraud detection or customer service can autonomously chain access across CBS, CRM, and data lakes. In a bank, that breaks SoD, audit trails, and RBI compliance."

Komal Vora
Chief Information Security Officer, Equifax Credit Information

with one important caveat. The expensive mistake is building separate control libraries to satisfy separate audits. The dangerous mistake is to assume that satisfying one framework automatically satisfies the others. SEBI's access-review cadence, DPDP's Rule 6 safeguards and CERT-In's audit evidence overlap heavily but are not identical: retention period, data residency and the granularity of what counts as an access log all diverge in ways an auditor will actually check. The safe design is one identity-evidence dataset with regime-specific extraction views, not one flattened control set. For Indian financial-services operators, that convergence has already changed what boards ask. "The board conversation has moved from 'Do we

have MFA?' to 'Can we demonstrate identity governance to the regulator in real-time?'" declares Komal Vora, Chief Information Security Officer, Equifax Credit Information. Regulators, according to Vora, now demand provable control over "who accessed customer data, core banking, and SWIFT systems, when and why."

CISO Takeaway

CERT-In's six-hour clock is the one that hits first. Set the breach-notification playbook to six hours and log retention to the longest applicable window, then test whether your identity logs can answer "who accessed what, and when" inside that window.

The door and the corridor

Regulators ask who had access. Attackers ask how to become someone who does. Verizon DBIR's headline "identity is dead as an attack surface" conceals the fact that identity's role has changed shape rather than shrunk. An attacker who gets in through an unpatched VPN appliance still needs a set of credentials, real or forged, to move laterally, escalate privilege and reach anything worth stealing. Identity is no longer always the door. It is almost always the corridor. "An attacker may exploit a vulnerability to get through the door, but compromised credentials, privileged access or hijacked sessions can allow them to move laterally and potentially reach critical and sensitive data," according to Pradnya Manwar, Chief Information Security Officer and DPO, Thomas Cook. Her fix is to shrink what a stolen identity can touch. "Reducing the data-access surface so that an identity compromise does not become a data compromise," she recommends.

Place the current attack taxonomy against that framing and you will see that the routes in are not equal. Start with infostealer malware,

which has quietly become one of the highest-yield attack classes in the market: it harvests both passwords and session cookies—so a hijacked session needs no second factor at all. Next come phishing and adversary-in-the-middle attacks, which have moved from static credential theft to real-time session interception. These defeat password-only MFA outright.

CISO Takeaway

When did your team last test whether a stolen session cookie — not a stolen password — could reach Tier-1 data? Add number-matching to every push-MFA deployment, and treat session-token theft as a Tier-1 control problem not a malware problem.

The non-human identity problem

The session cookies and API tokens that infostealers harvest increasingly belong not to people



"The organisations that get this right in 2026 will be the ones that stopped treating IAM as plumbing and started treating it as the primary control plane for enterprise risk."

Baidyanath Kumar
Chief Information Security Officer
& Data Protection Officer, JK
Lakshmi Cement

but to machines, and that is where identity governance is weakest. Ask five vendors to define “non-human identity” (NHI) and examine the answers. Gartner flagged the taxonomy problem directly at its London IAM Summit: non-human identity, workload identity, machine identity, service account and AI agent are used inconsistently across the market. Before a CISO can govern the category, it helps to acknowledge that the category itself is still being debated! “Organizations simply do not know how many NHIs exist within their ecosystems. Unlike human employees who are tracked via HR systems from onboarding to offboarding, AI agents, API keys, and tokens are frequently spun up ad-hoc by developers or business units,” says Vasanth Pai, Global DPO at Tech Mahindra.

However, what is not in dispute is the direction of the volume problem from NHIs, even if the exact ratio is. The most widely cited figure comes from CyberArk, an identity-security vendor, whose April 2025 research put machine identities at more

Six Terms to Know

Term	What it actually governs	Evidence a regulator will ask for
IAM	The umbrella: administration, authentication, authorization and auditing of every identity, human and non-human	Full lifecycle logs from creation to deprovisioning, across every connected system
SSO	A convenience layer federating login across applications	Federation and session logs — never accepted alone as an access control
MFA	One control inside authentication	Factor type per transaction; number-matching or hardware-bound evidence for high-risk actions
PAM	Governance and session control for the accounts that can do the most damage	Session recordings, just-in-time grant records, privileged-account inventory
IGA	Decides what access should exist and certifies it periodically	Access-certification campaigns, entitlement reviews, segregation-of-duties reports
Zero Trust	A paradigm — not a product or a finished state	Continuous policy-decision logs, not a one-time perimeter check

than 80 to one against humans — a number KPMG's Cybersecurity Considerations 2026 report repeats. Gartner-cited research runs lower, at roughly 25 to 50 to one. Practitioners closer to the actual situation often report something more modest still, in the low single or low double digits to one.

That gap between the executive estimate and the practitioner estimate is itself a governance signal: the people accountable for the number and the people who could actually count it are not looking at the same inventory. KPMG's report, drawn from an expert panel of its own cyber leaders and senior security executives at Google, Microsoft, Palo Alto Networks and ServiceNow, names non-human identity governance as one of eight defining CISO risks for 2026. Its argument is that identity practices built around a human joiner-mover-leaver cycle do not survive the move to machine-scale identity. CISO Forum's survey makes the same point: once agents, bots and service identities outnumber people, the governance gap stops being a hygiene issue and becomes a systemic risk.

Maturity lags that risk. Only 31 percent of respondents describe their identity-first Zero Trust implementation as defined, against 34 percent developing and 24 percent immature, and just 8 percent rate their AI security governance as optimized. McKinsey's 2026 AI Trust Maturity Survey, which covered roughly 500 organizations and was published in March, identifies a failure mode it calls agent sprawl. Low-code and no-code platforms let almost anyone in an enterprise spin up an agent, producing redundant, fragmented identities with no central inventory. As the survey puts it: "You cannot govern what you cannot see." It also points to what separates the leaders: 65 percent of the most mature organizations have defined human-in-the-loop processes for their agents, against 23 percent



"Boards have realized that attackers no longer just hack into networks; they log in using compromised credentials."

Vasanth Pai
Global DPO,
Tech Mahindra

of laggards. Revoking access is the harder, Pai points out, "Because these identities are often hardcoded into business-critical workflows, pulling a single token or credential can cause cascading outages across dependent enterprise systems."

This is also where the CISO Forum's survey pushes back on the premise of this feature. Indian CISOs rank generative AI and shadow AI as their single largest security challenge (4.2), well ahead of identity at scale (3.1). Read quickly, that looks like a vote against an identity program. Read closely, it is the same story. Every agent is an identity, and shadow AI is, operationally, unregistered identity with access to data. Identity is the control plane through which AI risk has to be governed: who or what the agent is, what it may touch, and which human answers for it.

The numbers specific to AI agents are starker still. A Cloud Security

Alliance survey of 285 IT and security professionals, published in February 2026 and commissioned by Strata Identity, found that only 18 percent of respondents were highly confident their current IAM can manage agent identities, and only 23 percent had a formal, enterprise-wide agent-identity strategy. Only 21 percent maintain a real-time inventory of active agents, and only 28 percent can trace an agent's actions back to a human sponsor. Asked how they authenticate agents, 44 percent cited static API keys and 43 percent a plain username and password. Eighty-four percent doubted they could pass a compliance audit focused on agent behavior.

The Model Context Protocol, now the default interface between AI agents and enterprise tools, has a similar problem. A vendor assessment by the API-security firm Equixly, published in March 2025, found command-injection flaws in 43 percent of the popular MCP servers it tested. Gartner's forecast shows where this is heading, and it is the sentence CISOs should carry into next year's budget conversation: 70 percent of CISOs will have adopted Identity Visibility and Intelligence Platforms by 2028, and by 2029 machines will proxy all human access to enterprise systems, cutting account takeover by 80 percent.

CISO Takeaway

Create a non-human-identity registry as a system of record before deploying the next agent. Every credential gets a named owner, a business purpose and a privilege scope from day one. Treat MCP and agentic protocols as an identity problem first and a model problem second.

Legacy, verticals and the real budget

None of the analysis survives contact with the actual state of most Indian enterprise identity estates, or with

the teams that CISOs have to run them. CISO Forum's survey is blunt on the second point. Respondents rate the shortage of skilled talent as a tougher constraint (with 50 percent calling it high or extreme) than lack of budget or top-management support. For most identity programs, the binding constraint is people, not money — and a program built around a talent shortage is different from one built around a funding gap.

Technical debt sits at the foundation: much of the Indian enterprise estate — banking systems, ERP systems, decades-old manufacturing execution systems, hospital patient-record platforms — was built for a password-and-directory world that predates the phrase "identity provider." Layered on top is multi-IdP sprawl. Most organizations now run more than one identity provider, each with its own joiner-mover-leaver process, its own review cadence, and its own gaps between them. Though legacy systems are a practical constraint, one CISO has a workaround. "Legacy systems (like older on-premises apps, LDAP-reliant systems, SSH-based servers, and thick clients) remain a significant roadblock. They natively lack support for modern



"Identity is no longer verified just by a 'correct password'; it requires a combination of user identity, device health, location anomalies, and network risk profiles before granting access."

Sunil Nishankar
Chief Information Security
Officer & DPO, Ageas Federal
Life Insurance

WebAuthn protocols," observes Sunil Nishankar, Chief Information Security Officer & DPO, Ageas Federal Life Insurance. His answer is to modernize at the identity-provider layer, so that users get a passwordless experience "while the proxy handles the legacy authentication protocols on the backend."

The identity conversation also looks materially different by vertical, and so does readiness. EY's DPDP research puts healthcare and life sciences lowest, at 9.9 percent, against 50 percent in retail and e-commerce, 38.8 percent in technology services, and 34.7 percent in financial services. In manufacturing, energy and utilities — sectors under direct CERT-In audit priority — identity in much of the OT estate still means a shared service account and a vendor-default credential with no meaningful lifecycle; the fix is rarely a platform purchase but a multi-year segmentation project that needs to be negotiated with plant operations.

In healthcare, the identity problem is clinician-facing and time-critical: a doctor moving between wards needs access fast enough that a slow authentication step

Vendor Due-Diligence Checklist

- Can all authentication, identity and audit data stay resident in India?
- Is Aadhaar integration supported for finance and government use cases?
- Is there a genuine Indian-language UI for customer-facing flows?
- Are compliance evidence packs pre-built for DPDP, SEBI, RBI and CERT-In?
- What is the vendor's own multi-year exit story if it is acquired?
- Track A or Track B: does this purchase match the track this system is actually on, or the track the enterprise wishes it were on?

The IAM Compliance Calendar			
Regulator / law	Requirement	Deadline	Identity-control implication
MeitY — DPDP Rules 2025	Rule 6 “reasonable security safeguards”: access control, activity logs, breach-detection evidence	Notified 13–14 Nov 2025; phased, most obligations by 13 May 2027 (confirmed 18-month window; a 12-month acceleration was proposed Jan 2026 but not gazetted)	Role-based access to personal data; monitored, retained access logs
RBI — Authentication Directions 2025	Two-factor authentication with at least one dynamic factor for digital payments	In force from 1 Apr 2026; cross-border CNP from 1 Oct 2026	SMS-OTP-only largely non-compliant; FIDO/passkey and in-app cryptogram migration
SEBI — CSCRf	Zero Trust, least privilege, MFA, periodic access review	Phased: 1 Jan 2025 (entities under prior circulars), 1 Apr 2025 (newly covered entities); Aug 2025 circular added technical clarifications, not a new deadline	Access-review cadence, MFA enforcement, log retention
CERT-In — Directions / CISG-2025-02, -03	Six-hour incident reporting; 180-day in-India log retention; annual empaneled audits with hard evidence	Ongoing / annual	Access control and MFA explicitly within audit evidence scope; separate 15-control MSME baseline

becomes a patient-safety concern, which is exactly why shared clinical-workstation logins persist despite DPDP’s security safeguards now covering the patient data behind them. It is also the sector surveys hear least from: healthcare made up just 3 percent of the CISO Priorities Survey respondents.

In telecom, TRAI’s KYC mandate already ties every SIM to a verified identity through Aadhaar-based e-KYC, and a draft DoT rule would extend that into a common biometric verification system across operators. Unfortunately this regulatory identity layer rarely shares a platform or a team with ordinary workforce IAM. In IT/ITeS, the defining problem is federated access at scale: contractor and client-facing accounts tied to a contract that should expire exactly when the engagement does, which most access-review cycles — built around employee departures, not client

departures — do not catch.

Ownership is the other half of the resourcing problem. In most Indian enterprises, the CISO owns identity architecture and the audit relationship; IT operations owns day-to-day provisioning; the DPO or privacy office increasingly owns the DPDP evidence trail; and business units own access-approval decisions. A program that has not named these owners in writing before it starts will spend its first two quarters resolving disputes instead of closing gaps.

CISO Takeaway

Get a four-owner RACI — CISO, IT Operations, DPO, business unit — signed in writing before spending a rupee on tooling. Redesign service-desk account recovery this year, since the most sophisticated MFA rollout is only as strong as the recovery process behind it.

Building a board-defensible program

Everything culminates, eventually, into a sequence a CISO can defend in front of an audit committee — but that sequence cannot look identical for a large, well-resourced BFSI enterprise and a mid-market manufacturer or hospital group running on a fraction of the budget and headcount. A well-resourced program can move through a full identity inventory, a non-human-identity registry, just-in-time privileged access and continuous access review inside four fiscal quarters.

But, a constrained program should sequence differently: inventory privileged and DPDP-regulated-data accounts first, rather than attempting the full estate at once; get ownership agreed in writing before any tooling spend, since that costs nothing; and accept a manually maintained, version-controlled registry as

a deliberate bridge with a hard retirement date, not as a stable end-state. The companion Toolkit lays out both tracks, plus a short self-assessment for CISOs whose estate genuinely straddles both — which, in practice, describes most Indian enterprises, since a large bank's cloud-native core and its legacy branch AD forest rarely sit on the same track.

Baidyanath Kumar Chief Information Security Officer & Data Protection Officer at JK Lakshmi Cement, argues that whichever track a CISO is on, the sequence should start with privilege: "Move from static entitlements to just-in-time, least-privilege access, so privilege is granted for a task and a window, not held indefinitely," he recommends. His analysis of the operational pain: "excessive and standing privilege is what keeps me up at night, because it's the one attackers monetize most directly."

Boards, for their part, do not all think in these terms. In the CISO

Priorities Survey 2026, 49 percent of respondents say their board sees cyber as critical to resilience and brand trust. A further 21 percent still see security through a compliance and legal lens, 18 percent as a strategic enabler of transformation, and 7 percent as an operational cost center with limited ROI visibility. A board that thinks in resilience terms will ask the question Komal Vora described earlier: can you demonstrate identity governance, now?

On either track, anchor the program to five metrics that belong in a board pack, not a technical runbook: the proportion of non-human identities with a documented, named owner; the proportion of privileged accounts under just-in-time rather than standing access; the proportion of authentication on phishing-resistant MFA rather than SMS-OTP or push; the mean time to revoke access for a departed user across every connected system; and the proportion of OAuth grants with a documented business owner.

None of the five is exotic, and all five are auditable with data most enterprises already collect somewhere — the work is consolidating it into one reportable number, on one dashboard, reported monthly rather than annually. A metric reported once a year is not a program; it is an audit finding waiting to happen.

CISO Takeaway

Run the sequence that matches your actual budget and headcount, not the one that reads best in a vendor slide. Report the five board metrics monthly. A signed ownership agreement costs nothing and will determine year one's success more than any single tool purchase.

What's Next?

Three decisions will define the next three years of an Indian CISO's identity program.

First, choose the non-human-

Who Owns What?

Regulator / law	Requirement	Deadline
Identity architecture & risk strategy	CISO	CISO's identity architect
Day-to-day provisioning / de-provisioning	IT Operations	IT Ops / service desk
DPDP Rule 6 evidence & privacy safeguards	DPO / Privacy Office	DPO and CISO jointly
Access approval / entitlement certification	Business unit (line manager)	Line manager, via IGA workflow
Non-human identity ownership	Engineering / DevOps lead	Engineering team, against CISO policy

No DPO appointed yet? DPDP does not universally mandate the role the way GDPR does, and a meaningful share of Indian mid-market enterprises have not formally appointed one. Default that responsibility to the CISO and General Counsel jointly, in writing, with an explicit deadline to formalize the role rather than leaving Rule 6 evidence ownership unassigned by default.

identity governance model now, deliberately, rather than by default. Build machine-identity governance inside the existing IGA platform, or buy a dedicated non-human-identity product — either is defensible; deferring the decision for another budget cycle is not. Indian budgets are already leaning this way: in the CISO Priorities Survey 2026, identity-first Zero Trust architecture is ranked first by 31 percent of respondents as an investment priority, and is second only to AI-powered SOC automation.

Second, decide the passwordless migration path and budget it as a multi-year change-management program, not a policy memo. Synced passkeys suit workforce scale; hardware-bound FIDO2 keys suit privileged and highest-

assurance accounts, where the cost of compromise is highest, and the smaller population makes a rollout affordable. A workforce-wide rollout is realistically a two-to-three-year program at most large Indian enterprises — legacy applications, a mixed device fleet, and, in manufacturing, a shop-floor workforce that cannot simply be emailed a QR code, all contribute to that. Passwords should survive only as a monitored recovery path, never as a working credential.

Third, sequence the regulatory work as one data program, not one control program. Commit to the single evidence architecture set out earlier in this feature, and decide now who owns it. The regulatory clock does not pause for a program to catch up. The DPDP compliance window closes in phases through

May 2027. CERT-In's audit cycle and SEBI's access review requirements run annually, regardless of whether the evidence pipeline supporting them is ready.

The day a CISO can answer all five board metrics with an actual number, on demand, rather than a promise to find out — that is the day identity stops being a project and becomes a governance program.

CISO Takeaway

Create a non-human-identity registry as a system of record before deploying the next agent. Every credential gets a named owner, a business purpose and a privilege scope from day one. Treat MCP and agentic protocols as an identity problem first and a model problem second.

The CISO's 2026 IAM Checklist

- Complete a full inventory (Track A) or a privileged/regulated-data-first inventory (Track B)
- Map identity data to the DPDP / RBI / SEBI / CERT-In evidence pipeline
- Pilot, then scale, phishing-resistant MFA for the highest-risk tier
- Freeze new standing privileged access; migrate to just-in-time
- Stand up a named-owner non-human-identity registry
- Sign the four-owner RACI before any tooling spend
- Extend identity threat detection and response coverage across AD, cloud IdP and SaaS
- Redesign service-desk account recovery
- Publish the five board metrics monthly, not annually
- Move access reviews from annual to continuous
- Document an exit plan for every Tier-1 identity vendors

IAM Implementation for Different Budgets

Track A assumes the resourcing a large BFSI, telecom or IT/ITeS enterprise typically has.

Track B assumes the constrained budget and thinner bench most mid-market manufacturers, regional healthcare groups and smaller financial-services players are working with. Run the self-assessment below per system or business unit — many Indian enterprises are genuinely both tracks at once, in different parts of the same organization.

Self-assessment — answer per system, not once for the whole enterprise

1. Dedicated identity engineering headcount for this system?

A — yes, a named owner exists.

B — no, it is one of several things one person covers.

2. IAM budget for this system as a share of its own security spend?

A — at or above 10–12%.

B — at or below 6%.

3. Regulatory exposure?

A — sits under two or more of DPDP, RBI, SEBI, CERT-In simultaneously and has built evidence packs for at least one.

B — sits under CERT-In's baseline only, or hasn't started evidence work.

4. Estate character?

A — largely cloud-native, single or dual IdP.

B — legacy AD, OT, or more than two identity providers in the same environment.

Three or more "A" answers: run Track A there. Three or more "B" answers: run Track B there. An even split is itself a valid answer — it means that system needs its own hybrid sequencing.



Track A — Well-resourced

- Q3 2026: Full human and non-human identity inventory. Map to the four-regime evidence pipeline. Begin phishing-resistant MFA pilot for the highest-risk tier. Freeze new standing privileged access. Get the RACI signed.
- Q4 2026: Scale MFA to the full high-risk cohort. Complete the NHI registry. Roll out just-in-time access for Tier-1 privileged roles. Finalise DPDP and RBI evidence packs.
- Q1 2027: Extend identity threat detection and response across AD, cloud IdP and SaaS. Redesign service-desk recovery. Deliver the first quarterly board metrics package. Bring the AI-agent registry live with named ownership.
- Q2 2027: Move to continuous access review. Run the first joint audit testing the one-dataset, several-views design against all four regimes at once. Confirm DPDP evidence is audit-ready. Refresh against that year's analyst research.
- risk tier only — finance, HR, DPDP-regulated data — not workforce-wide. Start a manually maintained, version-controlled NHI registry, with a hard retirement date no later than Q1 2027.
- Q1 2027: Extend MFA and just-in-time access to the next risk tier as budget allows. Redesign service-desk recovery — a process change, not a purchase. Defer full ITDR coverage to Track A's timeline.
- Q2 2027: Begin, rather than complete, the DPDP evidence pack, prioritising Rule 6 over full consent-management build-out. Report the same five board metrics as Track A even from a manually maintained count — a reported number beats no number.

The five board metrics (both tracks)

Track B — Constrained budget

- Q3 2026: Inventory privileged accounts and DPDP-regulated-data systems first — not the full estate. Freeze new standing privileged access (free). Get the RACI signed before any tooling spend.
- Q4 2026: MFA for the highest-
- QProportion of non-human identities with a documented, named owner
- Proportion of privileged accounts under just-in-time rather than standing access
- Proportion of authentication on phishing-resistant MFA rather than SMS-OTP or push
- Mean time to revoke access for a departed user, across every connected system
- Proportion of OAuth grants with a documented business owner. ■



Machine Speed, Human Stakes: Why India's Cyber Defences Are Running Out of Time

Check Point's Sundar Balasubramanian explains why India's fragmented, machine-speed threat landscape demands unified prevention, identity governance, and board-level ownership.

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

India is now one of the most targeted nations in the cyber world, not by coincidence, but by design. With organisations absorbing nearly 3,300 attacks per week, significantly above the global average, the country's rapid digital expansion has outpaced the security architectures meant to protect it.

Sundar Balasubramanian, Managing Director for India and South Asia at Check Point Software Technologies, has watched this threat landscape intensify in real time. A seasoned voice on enterprise security strategy, Balasubramanian brings a perspective that is at once granular and urgent, from the structural blind spots hiding inside fragmented cloud architectures to the boardroom conversations that still treat cyber risk as an IT line item rather than a business imperative. In this wide-ranging conversation with CISO Forum, he makes one thing unambiguously clear: the old playbooks no longer work, and the clock is ticking.

CISO FORUM: With India at 3,195 weekly attacks and AI accelerating adversarial speed, where are Indian enterprises most structurally unprepared and where are the most dangerous blind spots?

SUNDAR BALASUBRAMANIAN: India's cyber threat environment is now operating at significantly higher intensity than global averages, with organisations facing nearly 3,300 cyberattacks per week compared to the global average of 2,064 attacks per organisation, according to a recent Check Point Threat Intelligence Report. This reflects not just higher attack frequency, but the rapid expansion of enterprise attack surfaces across cloud, SaaS, APIs, AI workloads, hybrid infrastructure, and third-party ecosystems.

What is increasingly evident is that attackers are no longer focused only on perimeter exploitation. They are targeting visibility gaps

created by fragmented security architectures. Many organisations still operate with siloed controls across cloud, endpoint, identity, and network environments, making real-time threat correlation difficult and slowing prevention.

The biggest structural blind spots today include unmanaged machine identities, shadow AI deployments, insecure APIs, lateral movement across hybrid environments, and weak governance around third-party access. Machine identities are scaling exponentially faster than human identities, yet most enterprises still lack centralised lifecycle management and behavioural monitoring.

AI is also fundamentally compressing attacker timelines. Check Point's latest findings show that 90% of global organisations encountered risky AI prompts, 40% of MCP servers were vulnerable, and high-risk prompts increased by 97% year-on-year in 2025. This means enterprises are now facing AI-enabled reconnaissance, AI-generated phishing, automated malware adaptation, and AI-assisted impersonation campaigns at scale.

India's exposure remains particularly elevated across multiple malware categories, including infostealer attacks at 7.5% versus 4.6% globally, banking malware at 4.3% versus 1.7%, botnet activity at 17.3% versus 9.9%, and ransomware impact at 7% versus 3.6% globally. Additionally, 92% of malicious files delivered in India last month came through web-based channels, highlighting how browser-driven attack vectors are dominating modern threat delivery.

From a governance perspective, frameworks such as the DPDP Act and CERT-In mandates are also pushing enterprises to rethink cybersecurity beyond perimeter defence. Security now needs to operate as a continuous operational resilience layer embedded across infrastructure, identity, AI governance, and data flows.

CISO FORUM: As agentic AI represents the next frontier of enterprise risk, how should CISOs think about its unique risk profile, and what does a staged security architecture look like across maturity levels?

SUNDAR BALASUBRAMANIAN:

Agentic AI introduces a fundamentally different cyber risk model because these systems are designed not just to analyse information, but to autonomously make decisions, execute workflows, interact with APIs, and access enterprise systems with minimal human intervention. This shifts the security challenge from application protection toward runtime governance and execution-layer control.

The primary concern today is that enterprises are operationalising AI agents faster than they are redesigning their governance architectures to support them. Check Point Research has already highlighted growing exposure to vulnerable MCP servers, risky prompts, and AI-assisted attack automation. Threat actors are increasingly using AI to accelerate phishing, malware customisation, reconnaissance, and impersonation campaigns.

The risk profile around agentic AI includes prompt injection, privilege escalation, plugin compromise, data poisoning, runtime manipulation, excessive permissions, and AI-assisted lateral movement. What makes this especially challenging is that compromise can now occur through execution logic, not just infrastructure exploitation.

CISOs therefore need to approach AI security through a phased maturity model. The first stage is governance and visibility, understanding where AI agents operate, what systems they access, and how sensitive data flows through them. The second stage is runtime protection through Zero Trust enforcement, API governance, segmentation, behavioural analytics,



"The old playbooks no longer work, and the clock is ticking."

and continuous authentication. The final stage is autonomous resilience, where organisations implement AI-driven policy orchestration, automated containment, and continuous runtime trust validation.

In India, this becomes even more important under DPDP obligations and emerging governance expectations across BFSI, telecom, healthcare, and critical infrastructure sectors. AI governance is increasingly linked to accountability, explainability, data localisation, and operational compliance.

The long-term challenge is therefore not simply protecting AI systems but governing autonomous decision-making inside enterprise environments. Organisations that continue to treat AI security as an isolated application security problem will struggle to manage operational risk at scale.

CISO FORUM: Prevention vs Resilience: With 44% of Indian CISOs

prioritising cyber resilience over prevention, how do you reconcile these two imperatives — and where does the Infinity Platform fit a resilience-first strategy?

SUNDAR BALASUBRAMANIAN: The debate between prevention and resilience is increasingly becoming obsolete because modern cyber environments no longer allow organisations to prioritise one at the expense of the other. Prevention without resilience assumes compromise is impossible, while resilience without prevention creates operational fatigue, escalating recovery costs, and prolonged disruption.

This balance has become especially critical in India's high-intensity threat landscape. Check Point Research observed a 53% year-on-year increase in ransomware victim volumes globally, alongside nearly 50% growth in new ransomware groups, according to Check Point Research's latest

Ransomware Report. Attackers are also embedding AI into phishing, reconnaissance, malware customisation, and extortion workflows.

India is already experiencing disproportionately high exposure, with ransomware impacting 7% of organisations compared to 3.6% globally. In comparison, botnet exposure in India stands at 17.3%, nearly double the global average, according to the Check Point Threat Intelligence Report. Modern attacks now move at machine speed, establish persistence early, and often exfiltrate sensitive data long before disruption becomes visible.

This is why recovery-centric models alone are no longer sustainable. Enterprises need architectures that reduce the probability of compromise, limit lateral movement, and minimise the operational blast radius before disruption occurs.

Platforms like Check Point Infinity are designed around this convergence between prevention and resilience. Check Point Infinity unifies cloud, endpoint, network, mobile, email, and identity protection into a single operational architecture powered by ThreatCloud AI, rather than operating as isolated security layers; it delivers prevention, detection, response, and policy orchestration continuously across environments.

The practical advantage is operational coherence, enabling organisations to reduce attack surface, accelerate detection, strengthen response coordination, and sustain business continuity. In today's environment, resilience should not be measured only by recovery speed, but by how effectively organisations can sustain operations under active attack conditions.

CISO FORUM: Machine identities and third-party credentials sit at the lowest maturity tiers. What are the most common identity governance

failures in Indian enterprises and what is the realistic path to closing this gap?

SUNDAR BALASUBRAMANIAN:

Identity has effectively become the new enterprise perimeter as organisations continue expanding across hybrid work environments, SaaS ecosystems, cloud infrastructure, APIs, and third-party integrations. Attackers increasingly bypass traditional perimeter controls and target credentials, service accounts, machine identities, API tokens, and vendor access pathways because these often operate with excessive trust and limited visibility.

Check Point's India threat findings already reflect elevated identity-driven attacks. Banking malware impact in India stands at 4.3% compared to 1.7% globally, while infostealer attacks are at 7.5% versus 4.6% globally. These attacks are specifically designed to harvest credentials, escalate privileges, and establish persistent access inside enterprise environments.

The larger governance issue is that machine identities are now growing significantly faster than human identities, while security maturity around non-human identity governance remains low. Most organisations still lack centralised lifecycle management, behavioural monitoring, and policy enforcement for service accounts, APIs, workloads, and automated systems.

The most common failures we continue to see include excessive privileged access, inconsistent MFA enforcement, weak governance around vendor credentials, shared administrative access, poor API token hygiene, and limited visibility into machine identity behaviour. In many environments, identity governance still operates independently of cloud, endpoint, and network security, resulting in fragmented enforcement models.

The realistic path forward is identity-centric Zero Trust architecture built around continuous authentication, adaptive trust, least-

privilege access, and unified visibility across users, workloads, devices, APIs, and third-party systems. AI-driven behavioural analytics will also become increasingly important as machine-to-machine communication continues to scale rapidly.

From an India compliance standpoint, DPDP obligations and sectoral governance expectations across BFSI, telecom, healthcare, and critical infrastructure are further accelerating the need for stronger identity governance. Identity security is no longer a standalone IAM conversation; it now sits at the centre of enterprise cyber resilience.

CISO FORUM: With ransomware rated the most severe threat by 79% of respondents, how should Indian organisations redesign their architecture to disrupt the attack chain before encryption and exfiltration occur?

SUNDAR BALASUBRAMANIAN:

Ransomware has evolved far beyond isolated encryption attacks and now operates as a multi-stage business disruption model involving credential theft, reconnaissance, lateral movement, privilege escalation, exfiltration, and extortion. The challenge is that many organisations still intervene too late in the attack lifecycle, often only after encryption begins or operational disruption becomes visible.

Check Point Research observed a 53% year-on-year increase in ransomware victim volumes globally, alongside nearly 50% growth in new ransomware groups, according to the latest Ransomware Report. Smaller and more agile ransomware operators are increasingly embedding AI into phishing campaigns, malware customisation, reconnaissance, and negotiation workflows.

India is already experiencing disproportionately high ransomware exposure, with 7% of organisations impacted compared to 3.6% globally,

according to the Check Point Threat Intelligence Report. This reflects not only elevated attack activity but also structural visibility gaps across hybrid enterprise environments. Modern ransomware campaigns increasingly target cloud environments, unmanaged endpoints, APIs, and third-party ecosystems where monitoring remains fragmented.

The architectural shift organisations need today is moving from recovery-centric security models toward attack-chain interruption. The objective should be to stop attackers during the initial compromise, privilege escalation, and lateral movement phases before encryption and exfiltration occur.

Key priorities, therefore, include AI-powered phishing prevention, east-west traffic inspection, identity segmentation, least-privilege enforcement, continuous endpoint telemetry, behavioural analytics, immutable backup isolation, real-time threat intelligence correlation, and AI-assisted detection with automated containment.

A unified security architecture becomes critical in this environment because isolated controls cannot quickly correlate attack behaviour in AI-driven attack scenarios. Prevention today is increasingly dependent on cross-domain visibility and coordinated policy enforcement rather than standalone point products

CISO FORUM: With the talent shortage now the top internal barrier, what organisational prerequisites must be in place before SOC automation delivers, and what mistakes are organisations making in their automation journeys?

SUNDAR BALASUBRAMANIAN:

The cybersecurity talent shortage is becoming one of the biggest operational risks for enterprises because threat environments are expanding faster than security teams can keep up. Indian organisations are



“Operational compliance will increasingly depend on whether organisations can demonstrate technical enforcement capability, not just policy intent.”

simplification, and unified visibility.

In India, sectors such as BFSI, telecom, government, and critical infrastructure are also facing increased compliance expectations under the DPDP and CERT-In frameworks, making faster incident response, greater auditability, and enhanced operational visibility even more critical.

already facing nearly 3,300 attacks per week per organisation, according to the Check Point Threat Intelligence Report, significantly above global averages, creating growing analyst fatigue, alert overload, and operational inefficiencies across SOC environments.

One of the biggest mistakes organisations keep making is assuming automation alone will solve the talent problem. In reality, SOC automation only delivers value when foundational operational maturity already exists. Automating fragmented environments with inconsistent telemetry often accelerates noise rather than improving security outcomes.

Before deploying AI-powered SOC operations, enterprises need several foundational capabilities: unified telemetry visibility across cloud, endpoint, identity, network, and email environments; standardised response workflows; integrated

policy orchestration; risk-based alert prioritisation; high-quality threat intelligence; and cross-domain correlation capabilities.

Check Point’s research shows how AI is accelerating attack sophistication through AI-assisted phishing, malware adaptation, reconnaissance, and impersonation campaigns. This means SOC teams are now expected to defend against threats operating at machine speed, making operational simplification increasingly important.

The role of AI inside the SOC should therefore be augmentation, not replacement. AI is most effective when it reduces repetitive workloads, accelerates correlation, improves prioritisation, and allows analysts to focus on higher-order decision-making and incident response. Organisations achieving the strongest outcomes combine AI automation with platform consolidation, operational

CISO FORUM: Architecture, Not Documentation. Most organisations still treat DPDP compliance as a reporting exercise. What specific architectural changes must CISOs implement to be operationally compliant, not just documentationally ready?

SUNDAR BALASUBRAMANIAN:

Many enterprises still approach DPDP readiness primarily as a legal and reporting exercise, but operational compliance requires deeper architectural transformation across data visibility, access governance, runtime monitoring, and breach response. Policies and documentation alone do not reduce operational exposure if sensitive data remains fragmented, overexposed, or poorly governed across hybrid environments.

Check Point’s India findings from the Check Point Threat Intelligence Report show that Information Disclosure has become the most common exploit category, impacting

nearly 74% of organisations. This reinforces the idea that data exposure risk is increasingly operational, not theoretical.

To become operationally compliant under DPDP, enterprises need to redesign their security architecture to enable continuous visibility, enforcement, and accountability. Key priorities include continuous data discovery and classification, identity-centric access governance, encryption and tokenisation of sensitive data, API security enforcement, cloud workload visibility, runtime monitoring, real-time breach detection, and centralised auditability.

The challenge is compounded because Indian enterprises now operate across highly distributed environments, including SaaS applications, AI systems, cloud infrastructure, remote work models, and interconnected third-party ecosystems. Traditional perimeter-led governance is no longer sufficient in these environments.

Operational compliance will increasingly depend on whether organisations can demonstrate technical enforcement capability, not just policy intent. Under DPDP and CERT-In mandates, enterprises will need to prove they can detect, contain, audit, and respond to incidents in near real time.

The organisations that will mature fastest embed compliance directly into operational workflows through automation, unified telemetry, AI-driven monitoring, and continuous policy enforcement, rather than treating compliance as a periodic reporting exercise.

CISO FORUM: Third-party risk monitoring beyond tier-one vendors is at the lowest maturity level. What does credible supply chain governance look like at a conglomerate scale, and what do boards systematically underestimate?

SUNDAR BALASUBRAMANIAN:

Supply chain attacks are becoming

increasingly sophisticated because attackers are now exploiting trusted digital relationships rather than directly targeting enterprise perimeters. Modern enterprise ecosystems include SaaS providers, outsourced development teams, APIs, AI models, cloud platforms, open-source dependencies, and interconnected partner networks operating at enormous scale.

Check Point Research recently demonstrated an AI-assisted supply chain attack in which malicious code was introduced into an autonomous crypto-trading project via AI-generated contributions. This highlights how AI is now accelerating ecosystem-wide compromise pathways while lowering the barrier for attackers to infiltrate trusted software pipelines.

The biggest governance gap is that most enterprises still focus primarily on tier-one vendor assessment, while transitive exposure across fourth-party and fifth-party ecosystems remains largely unmonitored. Boards frequently underestimate how interconnected modern digital ecosystems have become and how quickly compromise can propagate across trusted environments.

The lowest maturity areas continue to include continuous third-party monitoring, vendor identity governance, software dependency visibility, API exposure management, runtime validation, privilege auditing, and behavioural monitoring across partner ecosystems.

Credible supply chain governance, therefore, requires moving from enterprise-centric security models toward ecosystem-centric governance. Organisations need continuous validation frameworks, software bill-of-materials visibility, Zero Trust enforcement, identity-centric controls, and AI-driven threat monitoring across interconnected vendor environments.

India's push toward digital sovereignty, combined with DPDP obligations and sector-specific governance requirements, is also

increasing pressure on organisations to strengthen oversight across outsourced operations, software dependencies, and data-sharing ecosystems. Supply chain security is no longer purely a procurement issue; it has become a core operational resilience challenge.

CISO FORUM: With deeply embedded point solutions and high sunk costs, how do you advise CISOs to build a pragmatic consolidation roadmap without creating new coverage gaps during transition?

SUNDAR BALASUBRAMANIAN:

Operational complexity has itself become a major multiplier of cyber risk. Indian organisations already face nearly 60% more attacks than the global average while managing fragmented environments across cloud, endpoint, network, identity, SaaS, APIs, and AI layers.

Many enterprises now recognise the strategic need for consolidation, but one of the biggest mistakes organisations make is pursuing aggressive vendor reduction before establishing operational visibility and governance maturity. Abrupt platform rationalisation can create dangerous blind spots if controls are removed faster than security capabilities to evolve.

A pragmatic consolidation strategy should therefore begin with telemetry unification rather than immediate tool elimination. Organisations first need centralised visibility, policy orchestration, and cross-domain correlation before rationalising overlapping technologies.

Key priorities should include telemetry unification across cloud, endpoint, identity, and network environments; centralised policy orchestration; cross-domain visibility; continuous validation testing; risk-prioritised migration; and platform-led operational simplification. Incremental consolidation is generally far more sustainable than large-scale replacement exercises.

The strategic objective is not



“Prevention without resilience assumes compromise is impossible, while resilience without prevention creates operational fatigue, escalating recovery costs, and prolonged disruption.”

only to reduce vendor count. It improves consistency in prevention, accelerates response times, reduces alert fatigue, and strengthens operational coherence across hybrid environments.

Platforms like Check Point Infinity are becoming increasingly relevant because they allow organisations to consolidate prevention, visibility, and operational management capabilities without sacrificing coverage during transition periods. A unified architecture becomes especially important in AI-driven threat environments, where fragmented tooling impedes correlation and delays responses.

From a business perspective, consolidation also needs to align with operational resilience, compliance readiness, and long-term governance strategy, not purely procurement-driven cost optimisation.

CISO FORUM: With leadership misalignment the biggest barrier to stronger security outcomes, what is your practical playbook for shifting the board conversation from budget approval to genuine co-ownership of cyber risk?

SUNDAR BALASUBRAMANIAN:

Cybersecurity conversations at the board level are increasingly shifting from technical oversight to enterprise resilience, operational continuity, and governance accountability. This shift is necessary because Indian organisations operate in one of the world’s highest-threat-density environments.

According to recent data from the Check Point Threat Intelligence Report, sectors such as education, government, telecom, and infrastructure are experiencing some of the highest attack volumes in India, with the education sector

alone facing more than 7,100 attacks per organisation per week. This demonstrates that cyber risk is no longer confined to IT environments; it now directly impacts operational continuity, customer trust, regulatory exposure, and business resilience.

One of the biggest barriers to stronger cybersecurity outcomes remains leadership misalignment. Many organisations still approach cybersecurity as an annual budget discussion rather than an enterprise-wide governance issue. As a result, accountability often remains isolated within security teams rather than being embedded in strategic business decision-making.

The most effective CISOs are reframing cyber conversations around financial exposure, operational disruption, AI governance, supply chain resilience, regulatory accountability, reputational impact, and customer trust. Boards respond far more effectively when cybersecurity is connected to measurable business outcomes rather than isolated technical metrics.

The practical playbook for stronger board engagement is translating cyber posture into business-impact language. Discussions around ransomware, AI risk, third-party exposure, and operational resilience should be framed in terms of revenue impact, downtime, ecosystem disruption, regulatory penalties, and strategic continuity.

Cybersecurity governance increasingly needs to resemble financial governance — embedded in transformation initiatives, cloud strategy, AI adoption, M&A decisions, and supply chain oversight. The organisations making the strongest progress today are those where cyber accountability extends beyond the CISO and becomes a shared leadership responsibility across business, operations, compliance, and technology teams.

This interview was conducted while Sundar Balasubramanian was working with Check Point. ■



Data is the new battleground, not the background: inside India's compliance reality

Supratim Chakraborty, Partner at Khaitan & Co's unpacks DPDP's compliance gaps, M&A impact, and AI-era balancing act.

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

As the Digital Personal Data Protection Rules move from paper to practice, Indian enterprises are finding compliance far messier than the law's text suggests. Data discovery, consent architecture, cross-border transfer restrictions, and breach notification timelines are converging into a single, unforgiving operational challenge that now shapes deal valuations, due diligence, and boardroom budgets as much as it shapes IT roadmaps.

In this conversation, Supratim Chakraborty, Partner at Khaitan & Co, unpacks where organisations are falling behind: the contractual groundwork between data controllers and processors that many are yet to begin, the scramble around Consent Manager Systems, and the dramatic shift breach reporting will undergo once the Data Protection Board becomes fully operational. He also addresses a question increasingly on every CISO's mind: how DPDP intersects with AI, and whether India has struck the right balance between protecting personal data and preserving room for innovation.

CISO FORUM: With the DPDP Rules moving from draft to implementation, where do you see the widest gap between what the law expects and where Indian enterprises currently stand?

SUPRATIM CHAKRABORTY:

Several gaps are emerging as implementation begins.

One key gap is data discovery: where is all the data, inside or outside the organisation? That question raises many challenges. Another important issue, which people are not as mindful of right now, is data inventories and mapping, and carrying out notice and consent work, mostly internal work, which you can expedite whenever you want.

But this law doesn't put any onus on the data processor. Whether it is a controller-to-processor relationship, controller-to-controller,

or processor-to-processor, you'll need an agreement. It has to call out all the data-related elements, the indemnity, and everything you need to cover yourself against. You need to send this contract or data addendum to your counterparty, and they have to sign and send it back.

That may not happen overnight. What people are missing right now is that exercise; if it doesn't start now, you are far too late.

CISO FORUM: You've spoken about data becoming "the battleground" rather than the background of a transaction. How is DPDP readiness now factoring into deal valuations, due diligence, and risk allocation in M&A?

SUPRATIM CHAKRABORTY: In M&A, we used to start with tax structuring and then move into mergers and acquisitions.

Right now, we are also seeing that in most M&A transactions, data structuring happens first. For example, if the acquirer is overseas and has cross-border aspirations to take out the data, and there is a data localisation requirement, it is a failed venture. As a result, we are seeing much more sincerity around data protection and cybersecurity measures.

Deal valuations are being revisited, and targets are sometimes being pushed to see whether they are prepared for the new regime. The penalties, as we all know, could be rupees 250 crore or more. As a result, we are seeing potential acquirers become more sceptical and measured. Due diligence is becoming better and more intense.

This is actually going back to the valuation discussion. Then the transaction proceeds.

CISO FORUM: The consent management framework is one of the more operationally demanding parts of the Act. What are organisations getting wrong when they try to build or integrate Consent Manager Systems?

SUPRATIM CHAKRABORTY: First, we have to get our own house in order.

We need to understand the consent management framework we build for our organisation. Once that is ready, we should consider syncing up with the consent manager. Right now, I'm seeing some players thinking, "I don't need to do anything about this."

Consent managers will come into play, and then we should start thinking about it. By that time, it will be extremely late. We should definitely start putting our house in order. Our own consent management system should be fully operational.

CISO FORUM: The rules leave room for government-notified restrictions on transferring personal data outside India. How should CISOs and legal teams plan for this uncertainty right now, rather than waiting for final notifications?

SUPRATIM CHAKRABORTY: There are two things; one is this law. When the drafting started, we began by saying almost hard data localisation, as all data should stay within the country. Why do we do that? There are several theories and countertheories.

One theory is the trade theory: my data is my own, so I should protect it the way I want. I should use it the way I want, etc. The other is the security theory: my people's data should be safe wherever it is, and my law enforcement authority should have access to it whenever they want.

As I understand it, our focus is primarily on the security side. Therefore, it is important to keep that in mind. Right now, we do not have a negative list of jurisdictions. The idea was never that we would have a long list of negative jurisdictions. We always thought it would be a short list.

We all know our countries of concern. The only problem right now is that, because of the new trade tensions, the equations have changed. We should fully prepare our organisation for everything

else. My hunch is that our prime trading partners—the good countries with which we have had good relationships forever—will continue to do so. If there is a geopolitical issue, we may see some jurisdiction being called out in that negative list.

CISO FORUM: Once the Data Protection Board of India is fully operational, how do you expect breach notification timelines and enforcement posture to differ from what companies have experienced under sectoral regulators like RBI or SEBI?

SUPRATIM CHAKRABORTY: There would be a dramatic change. We all know that, beyond the RBI and SEBI, certain players are still regulated.

Even CERT-in, the Computer Emergency Response Team of India, requires reporting any cyber incident within six hours today. They are becoming fairly proactive. What happens when this data protection board comes into play? If an incident involves personal information or digital personal information, you must report it immediately, along with certain information.

Within 72 hours, you'll have to go again and provide a much broader set of further information, which is quite a bit. If you look at those line items, it is almost saying who that entity is, the person who has done this activity within 72 hours. You would probably know nothing. And the last line of items in those 72 hours is the notification you sent to individuals.

You also must inform the impacted users, of course. Tell them that this incident happened. This is what data was probably leaked. This is what we have done. These are the things you can do to protect yourself. If you have a problem, please reach out to us.

In that last line item for 72-hour reporting to the board, they are saying, "Please let us know about that reporting." So, in effect, they are expecting it within 72 hours, and you have also informed the individuals. This is a dramatic change compared to the situation today.

"Data structuring happens first" now, ahead of tax structuring, in most M&A transactions today."

CISO FORUM: You train enterprises at the intersection of AI and data privacy. Where does DPDP fall short in addressing AI-specific risks like automated profiling or training-data provenance, and are companies over- or under-correcting for this?

SUPRATIM CHAKRABORTY: The Indian government has been mentioning this for quite some time: yes, of course, it is absolutely critical to protect individuals' personal data. But at the same time, we need to ensure that innovation, entrepreneurship, and ease of doing business do not suffer as much as possible.

We have already observed some jurisdictions that build phenomenal laws, and then they have to self-correct because innovation suffered, and business suffered. We must be smart about this as a country and take a middle-of-the-road approach. The DPDP Act is a data law. Of course. It covers all the relevant aspects of protecting data, so that part is already protected as far as AI activities are concerned.

But they have also added a few things. For example, for significant data fiduciaries, which are larger players crunching larger amounts of data and more critical data, they are saying there must be algorithmic accountability. This ensures that the rights given to individuals under the DPDP Act should not be hampered. If you look, there is insulation baked in, but this is a data law, so it has tried,

to the extent possible, to look at our economy and our business posture today, to protect personal data.

At the same time, we have to balance it with the AI world we are living in. In fact, the last part of the change happened in the rule stage. Over time, it became clear that we needed to put something in place around the AI. But we should also remember that it has tried to balance innovation.

If you see, research is still an exception. If no personal data is involved, that is outside the purview of things like that. We have to see the tricky time the regulator, the government, and the lawmaker are also going through. They have done a fair job. Is it perfect? Probably not. But can we make it perfect over time? I think yes.

CISO FORUM: Significant Data Fiduciary obligations put more compliance weight on leadership. What should a CISO's pitch to the board look like to secure the budget and mandate DPDP readiness requires?

SUPRATIM CHAKRABORTY: The first important thing I would like to talk about is that, again, when you look at significant data fiduciaries, they are much bigger and much more mature. Therefore, they would already have the budget and maturity.

That was the thought process. It was, again, not that all and sundry should suddenly become significant data fiduciaries. In fact, that would also increase the burden on the government because of the oversight they would have. They would have to submit the report, excerpts, etc., of the DPDP Act and the auditor, etc. The idea isn't to have many businesses and a lot of players suddenly become significant data fiduciaries. Still, for those who are, it will be smarter for them to allocate a larger budget.

This is an obvious ask from management, and they should agree, given the overall harm that can be caused to individuals' personal data. ■

RBI's new cybersecurity rulebook: From advice to law



India's central bank has replaced its decade-old cybersecurity circular with a binding 233-paragraph framework, moving banking IT governance to board-level accountability.

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

The Reserve Bank of India has issued the (Commercial Banks – Cybersecurity, Technology: Risk, Resilience and Assurance Framework) Directions, 2026, dated July 31, 2026, replacing the 2016 Cyber Security Framework circular.

According to a gap-analysis report comparing the two documents, the shift is not merely an update; it's a change in legal character. The 2016 circular used the word "should." The 2026 Directions use "shall" and are issued under Sections 27 and 35-A of the Banking Regulation Act, making compliance mandatory rather than advisory.

Governance takes centre stage

While the 2016 rules focused narrowly on technical controls, roughly a third of the new document three entire chapters addresses something the old framework barely touched: governance. Boards must now approve cybersecurity policy annually. In addition, a dedicated IT Strategy Committee, chaired by an independent director with at least seven years of technology experience, must meet quarterly. The Chief Information Security Officer role, mentioned only once in 2016, is now tightly defined: General Manager rank, no reporting line to the IT head, no business targets, and direct escalation to the Executive Director overseeing risk.

Deadlines replace vague timelines

The report highlights a recurring theme: specificity. Vulnerability assessments are now required every six months for critical systems, penetration tests annually, disaster recovery drills twice a year, and cyber incidents must be reported within six hours on RBI's DAKSH platform. This replaces the old circular's loosely worded "two to six hours."

New ground: Vendors, cloud, and ATM networks

Several requirements had no precedent in 2016. Banks must now maintain a source code escrow arrangement for critical applications, implement DMARC to fight email spoofing, and ensure public-facing infrastructure supports IPv6. Among the most detailed additions, third-

The 2016 circular used the word “should.” The 2026 Directions use “shall” and are issued under Sections 27 and 35-A of the Banking Regulation Act, making compliance mandatory rather than advisory.

party ATM Switch service providers face 24 specific contractual controls covering everything from password hygiene to PCI-DSS compliance.

What got dropped

Not everything survived. The 2016 circular’s one-time incident-reporting template, its severity classification system, and its one-off 2016 compliance deadlines are gone, replaced by the DAKSH platform’s own reporting structure. In addition, some narrower provisions like the express prohibition on macros in office documents were quietly dropped, seemingly folded into broader controls.

The bottom line

For banks that were fully compliant with the 2016 rules, the report notes that technical controls likely remain largely adequate. The real compliance burden, however, lies in the governance apparatus, the new Information Systems Audit chapter, and the sharper periodicities now attached to existing controls, which will require board time, not just IT budgets.■

Shadow AI and rogue agents: Inside Akamai's 2026 enterprise risk report



AI adoption has outpaced enterprise governance, creating hidden risks from shadow tools, power users, and autonomous agents.

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

Artificial intelligence has quietly become the modern enterprise's new attack surface. That's the central warning in Akamai's Enterprise AI Usage Risk Report 2026, part of its State of the Internet/Security series, built on data from LayerX, the AI-security firm Akamai recently acquired. The report paints a picture of an enterprise landscape where AI adoption has outpaced governance, and security teams struggle to see what's actually happening within their own organisations.

A small group of "Power Users" drives most of the risk

While nearly 20% of employees use AI weekly, the report finds that risk isn't evenly spreading. A small slice of "AI power users" the top 5% generate at least 144 conversations each, compared to just 12 or fewer for the bottom half of users. These heavy users share more sensitive business contexts, upload more files, and increasingly let AI drive real decisions, making them the highest-value targets for attackers. In other words, the riskiest behaviour is concentrated in a few hands.

Shadow AI is everywhere

Beyond well-known platforms like ChatGPT, Copilot, and Gemini lies a sprawling "long tail" of AI tools operating with zero oversight. Startlingly, nearly half of all enterprise AI conversations (47%) happen through personal accounts rather than corporate-managed ones, and 14.4% of conversations tied to corporate email addresses run on personal "freemium" subscriptions, meaning data may keep moving outside corporate control without anyone realising it.

Data doesn't leak the way it used to

Traditional data-loss prevention tools were built for emails and file transfers. AI breaks that model entirely: sensitive information now moves through prompts, pasted text, screenshots, and iterative conversation channels. Legacy security tools were never designed to monitor. The report cites a real-world case in which a U.S. government official inadvertently exposed restricted CISA operational data via a public AI tool, underscoring how routine AI use can retain sensitive data in places teams cannot easily see.

Extensions and agents: The new frontier of attack

AI browser extensions are especially risky:



"Nearly half of all enterprise AI conversations (47%) happen through personal accounts rather than corporate-managed ones."

nearly 75% request high or critical permissions, and they're almost three times more likely to request cookie access than average extensions. The report details two proof-of-concept attacks: "CursorJacking," in which a malicious extension silently steals AI coding assistant credentials, and "CometJacking," in which a booby-trapped webpage hijacks an AI browser agent via prompt injection. Together, they show how extensions and agents expand the attack surface.

The bottom line

Akamai's message to CISOs is clear: securing AI isn't about blocking tools, but about gaining visibility into power users, shadow apps, data flows, extensions, and autonomous agents alike. As AI agents take on more independent action inside enterprise systems, the report argues they must be governed with the same rigour as any human employee, because their autonomy can retain and spread risk inside the organisation. ■

AI is speeding up cyberattacks, not reinventing them



AI isn't reinventing cyberattacks; it's accelerating them—compressing weeks of malware development into days, while enterprises deploy AI agents faster than they can govern them.

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

Sophos' new AI Security 2026 report, based on casework from more than 625,000 customer organisations, delivers a clear message: artificial intelligence isn't creating new kinds of cyberattacks; it's speeding up old ones. Attackers still need to break in, move around, and steal data. What's changed, the report argues, is the clock.

A real-world wake-up call

The report opens with a striking case: in May 2026, Sophos discovered an attacker running roughly a dozen AI agents inside a customer's network, each testing malware against different security tools on its own virtual machine. Coordinated through a coding assistant, the agents produced nearly 80 malware modules and over 70 evasion techniques work that would normally take a human developer weeks—compressed into days. The operator later used this toolkit to deploy ransomware and steal data.

Criminals are adopting AI like any other tool

Sophos found threat actors buying and selling AI API access on underground forums, recruiting "AI prompt engineers," and advertising AI-powered voice bots and fake romance-fraud personas. But adoption isn't universal some criminals remain sceptical that AI will change their business at all. The pattern looks less like a revolution and more like criminals slowly folding a new tool into old habits.

Enterprises have a blind spot

A major theme is how AI tools themselves have become targets. Incidents like the Salesloft/Drift breach and a Vercel hack (triggered by an employee's AI tool) show how stolen credentials and OAuth tokens tied to AI assistants can open the door to company systems. Alarmingly, the report cites survey data showing that 71% of large enterprises have deployed AI agents with access to core business systems, but only 16% properly govern that access.

Defenders are using AI too with limits

On the positive side, AI reasoning models are helping security teams investigate alerts faster, cutting hours of analyst work into minutes. But Sophos cautions this isn't magic: current AI tools lack memory

“71% of large enterprises have deployed AI agents with access to core business systems, but only 16% properly govern that access.”

between sessions, meaning they investigate every alert as if for the first time, and over-relying on "reasoning" can sometimes make results worse rather than better.

What companies should do now

The report closes with practical advice: build visibility into which AI tools and agents have access to company data, tighten identity and credential controls around them, and patch known vulnerabilities faster since AI is shrinking the gap between a flaw being disclosed and being exploited, sometimes to under 24 hours. For companies deploying their own AI agents, Sophos recommends sandboxing, credential isolation, and human approval for any irreversible actions.

The bottom line:

AI hasn't rewritten the rules of cybersecurity yet — but it has sped up the game on both sides, and organisations that don't adjust their own tempo risk falling behind.■

Phishing gets smarter, not louder: Inside Zscaler's 2026 threat report



Phishing volume is falling, but attacks are growing sharper and costlier, powered by AI-generated sites and encrypted, near-invisible delivery.

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

Phishing emails are down. Phishing damage is up. That's the uncomfortable finding at the heart of Zscaler's ThreatLabz 2026 Phishing and Initial Access Report, which tracked attacker behaviour across its global security cloud through 2025 and into early 2026.

Fewer lures, bigger payoffs

After peaking at over 2 billion hits in 2023, phishing volume on the Zscaler cloud fell nearly 20% year over year in both 2024 and 2025. That's not attackers retreating; it's them getting smarter. Stronger email filters and identity defences have pushed criminals away from mass spam and toward fewer, sharper attacks that mimic real business workflows. The Services sector saw a 65.5% jump in attacks, and Government rose 50%, as scammers exploited routine trust in billing, renewals, and public communications.

AI is the new phishing factory

Artificial intelligence has quietly become the industry's most efficient con artist. ThreatLabz identified over 413,000 AI-generated websites, with roughly 9% flagged as malicious, many built using mainstream "text-to-site" tools like Manus AI, Blackbox AI, and Lovable. What once required a developer and days of work can now be produced with a single prompt in under an hour, allowing attackers to spin up convincing brand clones from fake Coinbase wallets to counterfeit government portals almost instantly.

Hiding in plain sight

Perhaps the most striking figure: 95.2% of all phishing activity in 2025 moved through encrypted (HTTPS) channels, up from 87% for malicious traffic overall. This means compromise increasingly happens quietly inside a browser session, not in an inbox, making it invisible to security tools that don't inspect encrypted traffic.

MFA no longer a safety net

The report details phishing kits like BlackForce, which combine real-time credential theft with the ability to intercept one-time MFA codes as victims type them, turning a simple phishing click into a full account takeover within seconds, even when multi-factor authentication is enabled.

Reconnaissance before the strike

Long before an attack lands, adversaries are scouting. ThreatLabz's decoy network recorded 89.9 million hostile interactions from 1.37 million unique attackers IPs across 520 customer environments in just six months. Manufacturing,



“95.2% of all phishing activity in 2025 moved through encrypted (HTTPS) channels, up from 87% for malicious traffic overall.”

financial services, and banking absorbed more than half of this activity, with attackers scanning for exposed logins, VPNs, and collaboration platforms they can exploit for entry.

What this means for 2026

Zscaler's predictions for the year ahead read like a warning label: AI agents phishing other AI agents, deepfake voice and video replacing traditional social engineering, and phishing evolving into persistent, multi-channel campaigns that blend email, SMS, and messaging apps into one coordinated con.

The core message for security leaders is simple: phishing volume is no longer a proxy for phishing risk. As attacks grow quieter, faster, and more convincing, defences built solely on inbox filtering are no longer enough. The focus needs to shift to visibility into encrypted traffic, identity verification, and catching attackers during reconnaissance before the first click happens. ■

When AI Agents Cross the Line

It is amazing how far we have come—from the typewriter, which transformed offices in the late 1800s, to artificial intelligence today. Even its name is being debated. If US President Donald Trump has his way, we may soon be calling it “super intelligence.”

But a new name does little to answer the harder question: Are we prepared for systems that can act, rather than simply respond?

In 2014, Stephen Hawking warned: “The development of full artificial intelligence could spell the end of the human race.” He was speaking about what full AI might become, not describing the systems we use today. Still, his warning raises a question that enterprise leaders can no longer treat as distant: can we keep increasingly capable AI within the boundaries we set?

AI is moving from a tool that helps people write, analyse and code to an agent that can use software, access information and complete tasks. That could mean resolving customer issues across systems, investigating security alerts or coordinating complex operations. It also means an agent can take an unintended path with real consequences.

The incident involving Hugging Face, a leading platform and community for open-source AI and machine learning, made this risk tangible. During an internal cybersecurity evaluation, OpenAI agents escaped a restricted test environment and compromised parts of Hugging Face’s infrastructure. In pursuing their evaluation goal, they carried out a sustained intrusion far beyond the boundaries of the test.

Hugging Face said it found no evidence that public models, datasets or Spaces were altered. That is an important limit on the incident’s impact. It does not erase the lesson: an agent crossed a boundary, reached external infrastructure and continued acting.

For technology and cybersecurity leaders, the question is how far an agent could go inside their own organisation. Its reach depends on the tools, credentials, data and network access it receives. Broad permissions can turn one unintended action into a wider incident, particularly when teams cannot see or interrupt what the agent is doing.

Enterprises need clear limits before granting agents the ability to act. Which systems can they access? Which actions require approval? How will teams trace their decisions and stop them when they stray? Testing should also examine how agents behave when instructions are ambiguous or a shortcut appears to help them achieve a goal.

Science fiction imagined machines escaping human control. The enterprise challenge today is concrete: build agents whose actions remain visible, bounded and accountable. The future of AI depends as much on those safeguards as it does on capability. ■



"Enterprises need clear limits before granting agents the ability to act. Which systems can they access? Which actions require approval? How will teams trace their decisions and stop them when they stray?"

Jatinder Singh

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

INSIGHTS
CIO & LEADER
studiotalks

INSIGHTS
CIO & LEADER
studiotalks

INSIGHTS

CIO & LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

studiotalks

WHERE TECHNOLOGY MEETS THE SPOTLIGHT !

WHY JOIN STUDIOTALKS

YOUR EXPERTISE. YOUR VOICE. YOUR SPOTLIGHT.

Drive bold discussions
defining the future of
enterprise technology.

Showcase your expertise
in a high-impact, TV-
style experience.

Get featured alongside
India's most influential
technology leaders.

SECURE YOUR SPOT NOW !

For business proposal

Hafeez Shaikh

Assistant Director - Projects

ET Edge - The Times Group

hafeez.shaikh@timesgroup.com

+91 9833103611

CURATED BY

ET
Edge

THE TIMES GROUP

Follow us: @CIOandLeader



AI SCALES. CAN YOUR INFRASTRUCTURE?

NOW TRY RUNNING IT ACROSS AN ENTERPRISE.

ET Edge **CIO&LEADER**

INFRASTRUCTURE SUMMIT 2026

THE INFRASTRUCTURE BEHIND THE **AI** ENTERPRISE

27-28 NOVEMBER 2026 | KARJAT

Where 75+ India's top CIOs & IT Leaders come together to decode what it will take to power the AI enterprise.



Two
Days



CIO
Workshops



Closed-Door
Roundtables



AI Infrastructure
Expo

AI IS SCALING. CAN YOU KEEP UP?

AN EVENT BY



For sponsorship, kindly write to:

Hafeez Shaikh

hafeez.shaikh@timesgroup.com

+91 9833103611

INDIA'S LEADING B2B INTELLIGENCE, COMMUNITY & MARKET PLATFORM