CIO Association

AI SECURED.
SECURED AI.

Building Cyber-Resilient Enterprises in the Age of AI

Key Insights from the CIO Association - India Cyber Security Symposium 2026

HOST ORGANISER: CIO Association, Delhi Chapter
SYMPOSIUM THEME: AI Secured. Secured AI.
DATE: 04 July 2026

About This Paper

This paper synthesises the collective intelligence of the **India Cyber Security Symposium 2026**, convened by the CIO Association. It draws on the day's expert keynote presentations, the practitioner panel discussions, the wider industry deliberations across the symposium tracks, and the dedicated discussion on India's Digital Personal Data Protection Act (DPDPA). Rather than a chronological record of proceedings, it is written as a working reference: a distilled account of what enterprise technology and security leaders are seeing, what has changed, and what now demands their attention.

It is intended for CIOs, CISOs, their Boards, and the wider community of enterprise-risk, legal, and business stakeholders who together own cyber resilience. The central conviction that ran through the symposium, and what organises this paper - is deceptively simple: **enterprises must secure their AI even as they increasingly use AI to secure the enterprise.** Everything that follows is an attempt to make that dual mandate concrete, measurable, and actionable.

Executive Summary

The India Cyber Security Symposium 2026 opened on a paradox that now defines the enterprise security posture: the same artificial intelligence that is transforming how organisations detect and contain threats is, in the same moment, transforming how those organisations are attacked. Security perimeters are mutating faster than governance frameworks can be written, reviewed, and codified - and 2026 is the year in which machine-speed adversaries moved from forecast to observed reality.

Three findings sit at the centre of the symposium's deliberations. First, the threat landscape has shifted qualitatively: autonomous AI agents can now map an enterprise's architecture, move laterally across segmented networks, and execute exploits with minimal human direction, compressing the defender's window from days to seconds. Second, AI systems have themselves become high-value targets - the data pipelines, training sets, and inference interfaces behind automated decisions are a first-class attack surface that most enterprises still under-protect. Third, cyber risk has become an economic and macro-stability concern of a scale that places it unambiguously on the Board agenda, reinforced by a regulatory environment in which the DPDPA attaches direct, personal accountability to leadership.

From these findings, the symposium converged on a set of recommendations for the year ahead. Enterprises should govern AI and agentic systems with the same rigour applied to any other production system, keeping every autonomous action within pre-approved, logged, and reversible boundaries. They should treat identity - human, privileged, and increasingly non-human machine identity, as the operative security perimeter. They should extend assurance beyond their own walls to third parties and the supply chain, and they should shift the strategic objective from cyber defence to cyber resilience: the demonstrated ability to absorb a breach and restore critical services at speed. Underpinning all of this, DPDPA readiness must move from policy on paper to operational reality, and progress must be made visible to the Board through a small set of meaningful KPIs and KRIs.

The remainder of this paper develops each of these themes, reconciles national-sovereign and enterprise perspectives on defence, and closes with a proposed CIO/CISO agenda for 2026–27, ten headline takeaways, and a clear call to action. **The message is consistent throughout: secure the AI you build, and use AI to secure the enterprise you run.**

1. What Changed in 2026 : The Emerging AI-Led Threat Landscape

If a single question framed the symposium, it was this: what is materially different about 2026? The answer is that the adversary now operates at machine speed. Where automated tooling once accelerated human attackers, autonomous agents now compress reconnaissance, intrusion, and exploitation into a continuous, self-directed sequence. The defensive implication is structural rather than incremental and the assumption of a human-paced attack, around which most incident-response playbooks are still written, no longer holds.

The Paradox of Velocity

India occupies a distinct, high-stakes position within the global technology architecture. The nation has established unprecedented digital scale, serving as an expansive engine of user activity, core infrastructure deployment, and cross-border software integration. However, this accelerated digital adoption has simultaneously widened the aggregate domestic attack surface, drawing the focus of sophisticated international adversaries. The contemporary enterprise is caught in a structural mismatch: security perimeters are mutating faster than legacy corporate governance frameworks can be written, reviewed, and codified.

Documented Threat Benchmarks

The discussions moved past theoretical abstractions by grounding operational realities in definitive, recent intrusion profiles:

- **Sovereign surveillance infrastructure:** Organised operations by state-aligned threat actors - specifically the sustained campaigns orchestrated by the group known as APT36, have intentionally bypassed standard enterprise barriers. These operations directly target strategic state institutions, leading academic research universities, and core public utilities with the explicit operational directive of maintaining long-term, undetectable network persistence and continuous intelligence extraction.
- **Exfiltration mechanics:** The commercial manufacturing sector experienced acute vulnerability, highlighted by the "World Leaks" ransomware operation. This single campaign resulted in the public exposure of over 200,000 internal, proprietary corporate files belonging to a large Indian electronics firm on dark-web repositories, emphasising that physical supply-chain leaders are now primary targets for economic extortion.
- **Autonomous infiltration vectors:** Network defenders are now confronting the deployment of fully autonomous AI coding agents by malicious actors. Documented incidents demonstrate that these threat models can independently map target architectures, orchestrate lateral movements across segmented networks, and execute tactical exploit payloads with minimal to no human command inputs.

2. Securing Enterprise AI and Agentic AI

The symposium's theme - AI Secured, Secured AI, captures a duality that enterprises can no longer treat as two separate conversations. AI is simultaneously the most powerful addition to the defender's arsenal and one of the most consequential new attack surfaces on the estate. The sections below hold both truths together: AI as a defensive multiplier, AI as an asset that must itself be assured, and the distinct governance burden created by agentic systems that act, and not merely advise.

AI for Security : The Defensive Multiplier

The integration of machine learning into modern Security Operations Centres (SOCs) marks a structural shift from human-dependent triage to machine-speed response.

- **Operational scale:** Traditional SOC models are fundamentally incapable of parsing the sheer volume of telemetry generated by modern cloud-native enterprises. AI is currently utilised to dynamically map baseline network behaviour, synthesise disjointed alerts into coherent threat chains, and isolate suspected nodes before a human analyst can review the initial telemetry.
- **The velocity paradigm:** Real-time AI threat detection transforms the standard defence playbook from a historical, retrospective posture ("What was breached?") to an active, predictive capability that counters machine-speed automated adversaries on equal footing.

Security for AI : Assured Intelligence

As machine-learning components migrate out of isolated testing sandboxes and embed into core corporate decision layers, the models themselves become high-value targets.

- **The unchecked attack surface:** Enterprise technology teams frequently underestimate the vulnerabilities inherent in their data pipelines, training datasets, and inference APIs. Adversarial manipulation, including data injection, model poisoning, and prompt extraction can compromise the integrity of automated decision engines.
- **Sovereign and corporate sanctity:** Securing the pipeline is an existential requirement. If the underlying dataset or model logic is quietly corrupted, the subsequent operational actions taken by the enterprise will be flawed from inception, making standard endpoint protection irrelevant.

The New Frontier : Governing Agentic AI

The most significant escalation discussed at the symposium was the enterprise adoption of agentic AI: systems that do not merely generate recommendations but take actions - querying systems, moving data, initiating transactions, and orchestrating other tools, all with a degree of autonomy. Agentic AI collapses the traditional gap between decision and execution, and in doing so it inherits, and amplifies, every weakness in the identity, data, and access layers beneath it.

- **The action boundary:** Every autonomous agent should operate inside an explicitly defined, pre-approved envelope of permitted actions. Where an agent can read, write, transact, or trigger downstream systems, those permissions must be scoped to least privilege, logged in full, and critically, be reversible.
- **Identity for non-human actors:** Agents are, in effect, a new class of privileged user. They require their own strongly authenticated, individually attributable machine identities, not shared service credentials, so that every action can be traced to a specific agent and revoked without collateral disruption.
- **Prompt and instruction integrity:** Because agents act on the instructions they are given, prompt injection and instruction manipulation become paths to real-world consequence rather than merely erroneous text. Guarding the instruction channel is now a security control, not a quality concern.
- **Human authority at the threshold:** The consensus was firm: the more consequential the action, the closer the human must remain. Autonomy is welcome for routine, high-volume, low-blast-radius decisions; it must stop at the threshold of irreversible or enterprise-critical action.

3. Domain Deep Dives and Expert Insights

Operational Technology (OT) and Plant-Floor Convergence

The integration of physical manufacturing plants with enterprise IT networks has created critical vulnerabilities, particularly within highly regulated spaces like pharmaceutical and automotive manufacturing.

- **The legacy deficit:** Industrial control systems (ICS) and plant-floor operational technologies were fundamentally designed for long-term mechanical reliability, completely isolated from external networks. Connecting these environments to cloud-based optimisation platforms exposes legacy equipment to automated external threats.
- **The threat profile:** In environments containing sensitive intellectual property, proprietary formulations, and automated manufacturing steps, an AI-accelerated incident bypasses standard human response timelines. The consensus shows that plant-floor defence models must transition away from legacy periodic updates toward continuous, isolated behavioural monitoring capable of containing anomalies instantly.

User Awareness Alone Is No Longer Sufficient

The symposium was careful to reframe, not dismiss, the role of the human layer. End-user security awareness remains a necessary foundation, but it can no longer be the load-bearing pillar of enterprise defence. Automated, AI-crafted social engineering has advanced to the point where even well-trained, attentive employees cannot reliably distinguish a malicious message from a legitimate one. The correct conclusion is not that awareness fails, but that awareness alone is insufficient and must be paired with controls that assume the human perimeter will occasionally be breached.

- **Behavioural harvesting:** Malicious AI tools can harvest large volumes of public and corporate behavioural data to craft highly personalised phishing vectors. These models can convincingly mimic the exact tone, linguistic style, and functional context of legitimate internal communications.
- **Systemic pressure:** Because these communications blend seamlessly with daily corporate workflows, they regularly bypass both standard email filtering and human scrutiny. Strategy must therefore complement education with real-time behavioural containment and strict zero-trust credential verification - so that a single convincing message cannot, by itself, become a breach.

The Bounds of Autonomy : Human-in-the-Loop Realities

As defensive AI systems gain the capability to autonomously isolate threats, technology leaders face difficult decisions about where to draw the boundary of autonomous execution.

- **The isolation trade-off:** While automated systems can instantly block specific IP addresses, isolate individual endpoints, or revoke compromised user privileges, granting them the authority to take wider actions presents operational risk.
- **The boundary rule:** The consensus among senior practitioners dictates that critical remediation actions such as shutting down production plant networks, cutting off primary database connections, or taking core business-continuity platforms offline - must be kept behind a strict human authorisation barrier. Defensive autonomy must be balanced against the risk of an automated system accidentally triggering a massive self-inflicted denial-of-service event.

4. Identity as the New Security Perimeter

With the network boundary effectively dissolved by cloud, mobility, remote work, and now autonomous agents, the symposium repeatedly returned to a single organising principle: identity is the new perimeter. In an environment where the human edge can be socially engineered and the network edge no longer contains anything, the credential and the trust decision made each time it is used becomes the true line of defence.

- **Three populations, one discipline:** Modern enterprises must govern workforce identities, privileged administrative identities, and a fast-growing population of non-human and machine identities - service accounts, API keys, workloads, and AI agents. The last of these is now often the largest and least governed.
- **Phishing-resistant, just-in-time access:** Static passwords and standing privileges are the raw material of most modern intrusions. The direction of travel is phishing-resistant multi-factor authentication for people and strong cryptographic identity for machines, combined with just-in-time, least-privilege access that grants entitlements only when needed and withdraws them automatically.
- **Continuous verification over one-time trust:** A zero-trust posture treats every access request as unverified until proven otherwise, evaluating identity, device, and context continuously rather than at a single gate. This is the practical mechanism by which an enterprise limits the blast radius of an inevitable compromise.

5. Third-Party and Supply-Chain Cyber Risk

A recurring theme across the panels was that an enterprise's security is now only as strong as the weakest link in its extended ecosystem. The most damaging incidents increasingly enter not through the front door but through a trusted supplier, a managed-service provider, a software dependency, or an integration partner. The "World Leaks" exposure discussed earlier is a case in point: a compromise at a supply-chain participant becomes an economic and reputational event for everyone connected to it.

- **Transitive trust is transitive risk:** Every vendor with access to systems or data extends the attack surface. Fourth-party dependencies, the suppliers of your suppliers, concentrate risk in ways most enterprises have never mapped.
- **From point-in-time to continuous assurance:** An annual questionnaire is no defence against a machine-speed adversary. Leading practice is shifting toward continuous monitoring of critical vendors, contractual security and breach-notification obligations, and a maintained inventory of who can touch what.
- **Software supply chain and AI provenance:** The provenance of software components, models, and training data is now a security property. Enterprises must be able to answer, for any critical system, where its components and its models came from and whether they can still be trusted.

6. From Cyber Defence to Cyber Resilience

Perhaps the most important shift in mindset urged at the symposium was the move from a defence posture to a resilience posture. Defence asks whether an attacker can be kept out; resilience accepts that, given a sufficiently determined and well-resourced adversary operating at machine speed, some intrusions will succeed, and asks instead how quickly and how completely the enterprise can absorb the impact and recover. In a landscape of inevitability, the organisations that endure will not be those that are never breached, but those that are never brought down.

- **Design for containment and recovery:** Resilience is engineered before an incident, not improvised during one. It depends on segmentation that limits lateral movement, on tested and isolated recovery paths, and on immutable backups that an attacker cannot reach or corrupt.
- **Recovery as a measured capability:** Recovery-time and recovery-point objectives (RTO/RPO) for tier-1 services should be defined, rehearsed through live restoration tests, and reported to the Board. An untested recovery plan is a hypothesis, not a control.

- **Resilience is organisational, not only technical:** The capacity to keep operating spans crisis communication, legal and regulatory response, customer trust, and business continuity. Cyber resilience is therefore an enterprise competency owned across functions, not a task delegated to the SOC.

7. The Economics of Cyber Risk and the Board Agenda

The Economic Cost of Frictionless Crime

The financial impact of modern cyber operations on the domestic economy has reached unprecedented levels. Data shows that economic losses from cybercrime across India reached Rs 22,495 crore in 2025 alone, driven by more than 28 lakh formalised law-enforcement complaints. This massive financial drain underscores that cybersecurity is no longer just a technical operational cost, but a primary threat to macroeconomic stability and corporate valuation.

Cybersecurity as a Board and Enterprise-Risk Agenda

Losses of this magnitude reframe cybersecurity from a technology function into a matter of enterprise risk, corporate valuation, and fiduciary duty. The symposium was unambiguous: cyber risk belongs on the Board agenda as a standing item, discussed in the language of business impact rather than technical controls. Boards do not need to become security experts; they need to be able to ask the right questions and to see a small number of meaningful measures that tell them whether exposure is rising or falling.

- **A business conversation, not a technical briefing:** The Board's interest is in the concentration of risk, the adequacy of resilience, the exposure created by third parties, and the enterprise's readiness to meet its legal obligations, expressed as business outcomes, not tool inventories.
- **Accountability that reaches the top:** Under the DPDPA and the broader governance environment, accountability for cyber and data outcomes now extends to the leadership and the Board. Oversight is a duty, not a delegation.
- **Measured, not asserted:** Confidence should rest on evidence. The KPIs and KRIs proposed later in this paper are offered as a starting point for a Board-level view of progress and exposure.

8. DPDPA : From Statute to Enterprise Readiness

The implementation of rigorous data-privacy frameworks, specifically India's Digital Personal Data Protection Act (DPDPA), has fundamentally shifted corporate accountability. Enterprise technology leaders now face direct civil and, in defined circumstances, personal liability for architectural oversights. Organisations can no longer treat technical procurement and legal compliance as separate business functions; modern enterprise defence requires an integrated strategy spanning technical prevention, real-time containment, strict regulatory alignment, and legal-defence readiness.

The dedicated DPDPA discussion at the symposium pressed beyond the letter of the law to the practical question every enterprise now faces: what does readiness look like? The answer offered was that DPDPA readiness is an operating capability, not a policy document, and that most organisations are further from it than their compliance paperwork suggests.

- **Know your data:** Readiness begins with a current map of personal-data processing - what is collected, on what lawful basis, where it flows, who can access it, and how long it is retained. An enterprise that cannot describe its data flows cannot protect them or account for them.

- **Operational consent, rights, and retention:** Consent capture, data-principal rights fulfilment, and retention limits must be working processes embedded in systems, not intentions in a policy. These are engineering commitments as much as legal ones.
- **Breach notification on the clock:** The Act's notification expectations assume an enterprise that can detect, assess, and report an incident within defined timelines. This links data-protection readiness directly to detection-and-response maturity discussed elsewhere in this paper.
- **Audit-ready evidence:** Because accountability is personal and demonstrable, the enterprise must be able to show its work - the records, decisions, and controls that evidence compliance if and when a regulator or an incident calls for them.

9. Reconciling National Sovereign and Enterprise Architectures

The symposium compared macro-level sovereign defence strategies with everyday commercial enterprise applications, highlighting how national-security principles translate into corporate IT architecture:

Strategic Domain	National Sovereign Framework (e.g., Israel / India)	Enterprise Corporate Architecture
Operational Mandate	Continuous state-level resilience against persistent, well-funded nation-state adversaries.	Protection of corporate intellectual property, operational uptime, and customer data registries.
Core Strategy	Unified national defence platforms, cross-border intelligence pooling, and protection of critical infrastructure.	Multi-layered enterprise security suites, automated SOC monitoring, and zero-trust network segmentation.
Regulatory Driver	National-security directives, defence doctrines, and international telemetry-sharing agreements.	Strict statutory data-privacy law (DPDPA), industry compliance metrics, and direct executive liability.
Data Architecture	Deeply isolated sovereign systems, national AI research missions, and secure defence communication channels.	Commercial cloud services, data pipelines, and public-facing APIs that require continuous monitoring.

10. The CIO / CISO Cybersecurity Agenda for 2026-27

Drawing the symposium's deliberations together, the following agenda sets out the priorities that enterprise technology and security leaders should carry into 2026–27, each expressed with an illustrative, Board-level KPI (a measure of progress) and KRI (a measure of exposure). These are offered as a practical starting point, deliberately few in number, to give Boards a defensible view of whether the enterprise is getting safer, and where its risk is concentrated.

Strategic Priority	Illustrative KPI - measure of progress	Illustrative KRI - measure of exposure
Secure-AI governance	% of production AI/ML systems with documented data lineage, a model inventory entry, and assurance sign-off.	Number of AI systems in production without a named owner or completed guardrail review.
Agentic AI guardrails	% of autonomous-agent actions executing within pre-approved, logged, reversible boundaries.	Count of privileged or irreversible actions taken by agents without a human authorisation checkpoint.

Strategic Priority	Illustrative KPI - measure of progress	Illustrative KRI - measure of exposure
Identity as perimeter	% of workforce, privileged, and machine identities under phishing-resistant MFA and just-in-time access.	Number of standing privileged credentials and orphaned or unattributed service accounts.
Third-party & supply chain	% of critical vendors under continuous monitoring with current security and breach-notification obligations.	Mean time to identify and contain a third-party-originated exposure; count of unmapped fourth-party dependencies.
DPDPA readiness	% of personal-data processing activities mapped, with consent, retention, and breach workflows operational.	Number of high-risk processing activities lacking a lawful-basis or impact record.
Detection & response velocity	Mean time to detect / respond (MTTD/MTTR) against machine-speed benchmarks; % of alerts auto-triaged.	Attacker dwell time on critical assets; backlog of un-triaged high-severity alerts.
Cyber resilience & recovery	Tier-1 RTO/RPO validated by live restoration tests within the reporting period.	% of tier-1 systems without a tested, isolated, immutable recovery path.
The human layer	Reporting rate (vs click rate) on simulated social-engineering tests; median time-to-report.	% of simulated intrusions reaching sensitive systems despite awareness controls.

11. Top Ten Takeaways from the Symposium

1. **The perimeter now mutates faster than governance can be written.** The operating assumption must shift from prevention alone to continuous validation and rapid containment.
2. **2026 is the year of the machine-speed adversary.** Autonomous AI agents can map, move, and exploit with minimal human input, compressing the defender's response window to seconds.
3. **Secure the AI you build, and use AI to secure the enterprise you run.** The data pipelines, training sets, and inference APIs behind decision engines are now first-class attack surfaces.
4. **Agentic AI raises the stakes.** Every autonomous action needs a pre-approved, logged, and reversible boundary, with critical remediation kept behind human authorisation.
5. **User awareness alone is no longer sufficient.** Training remains necessary, but defence must assume the human perimeter will be breached and invest in real-time behavioural containment and zero-trust verification.
6. **Identity has become the new perimeter.** Workforce, privileged, and non-human machine identities all require phishing-resistant authentication and just-in-time, least-privilege access.
7. **Third-party and supply-chain exposure is now a primary route of compromise.** Continuous vendor assurance matters as much as internal controls.
8. **Cyber risk is an economic and macro-stability issue.** With documented national losses at scale, cybersecurity belongs on the Board agenda as enterprise risk, not an IT line item.
9. **DPDPA turns data protection into personal accountability.** Readiness means mapped processing, operational consent and breach workflows, and audit-ready evidence — not policy on paper.

10. **Resilience is the real objective.** The enterprises that endure will be those that can absorb a breach and recover tier-1 services quickly, treating technology and legal-compliance readiness as a single, Board-owned priority.

12. Conclusion : A Call to Action

The India Cyber Security Symposium 2026 did not end on a warning, and neither should this paper. The threat landscape is real and it is accelerating, but the deliberations of the day pointed consistently toward agency rather than alarm: the tools, the practices, and the governance models required to meet this moment already exist, and the enterprises that adopt them deliberately will be the ones that thrive.

The call to action is the theme itself, now made operational. Secure your AI: govern your models, pipelines, and agents with the same rigour you apply to any critical system, and keep human authority at the threshold of consequential action. Use AI to secure the enterprise: let it carry the machine-speed detection and containment burden that humans no longer can. Make identity your perimeter, extend assurance to your supply chain, and measure resilience - the capacity to recover, as seriously as you measure defence.

For CIOs, CISOs, and their Boards, the mandate for 2026–27 reduces to a single, memorable line: secure AI while using AI to secure the enterprise. The organisations that internalise both halves of that sentence - and hold themselves to a small set of honest measures against it - will be the ones still standing, and still trusted, at the next symposium.

Contributor Registry

The insights, strategic frameworks, and tactical guidance in this compendium were synthesised from the presentations, panel debates, and strategic briefings delivered by the following experts, policymakers, and security practitioners:

- **Abhishek Bansal** — CISO, Axis Maxlife Insurance
- **Ajit Pillai** — Chief Business Officer, 63SATS Cybertech
- **Avaneesh Vats** — VP (Digital Infra & Innovation) and CIO, Techno Electric & Engineering Co. Ltd.
- **Barjesh Jain** — Senior Advisory Solution Consultant, ServiceNow
- **Basant Chaturvedi** — Delhi Chapter President, CIO Association
- **Deepak Rana** — Global CISO, Evalueserv
- **Dhananjay Khanna** — CISO, SBI Cards
- **Harsha Vardhan J** — Chief Information Security Officer & Data Protection Officer, Birlasoft
- **Himanshu Gautam** — Founder, GoTrust
- **Lt. General Madhavan Unnikrishnan Nair (Retd.)** — Former National Cyber Security Coordinator, Government of India; Signal Officer-in-Chief, Indian Army
- **Mahendra Manjunath** — Lead Industrial Consultant, BXI
- **Muneer KongaWani** — DGM & CISO, J&K Bank
- **Murtaza Bhatia** — Vice President & Practice Head – Cyber Security, Sify Technologies Ltd.
- **Nantha Ram Ramalingam** — Director & Leader – Cybersecurity, GCC, Albertsons Companies
- **Nisha A B** — Kerala State Financial Enterprises
- **Priyadarshi Achar** — Sales Engineering Leader, Proofpoint
- **Rajender Bhandari** — CTO, Binary

- **Sai Pavan Kumar Surapaneni** — CISO, Granules India
- **Sh. Narendra Nath Gangavarapu** — Joint Secretary, National Security Council Secretariat (NSCS)
- **Sonia Swami** — CISO, O2 Power
- **Srinivas L.** — Joint Managing Director and Joint CEO, 63SATS Cybertech
- **Umesh Mehta** — President, Governing Body, CIO Association
- **Unique Kumar** — Group CIO, CK Birla Group
- **Vinayak Ghodse** — CEO, Data Security Council of India (DSCI)
- **Vishwas Pitre** — CISO & DPO, Zensar Technologies
- **Yask** — Cyber and Privacy Professional, Indian Oil Corporation Limited (IOCL)
- **Yigal Unna** — Former Director General, Israeli National Cyber Directorate (INCD)