



RACE WEEKLY INTELLIGENCE BRIEF

RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER: **CISO Elevated: When Security Becomes the Board's Business**

ALSO INSIDE: **Deep Algorithms**
· **DPDP Deadline** · **Zero Trust Gap**

FROM THE EDITOR'S DESK

Issue #1 asked whether your organisation could inventory its AI agents. This week, we ask the harder question: once you find them, who is responsible for securing them? India's enterprises are deploying AI at speed while their security architecture was built for a world that no longer exists. DPDP enforcement is six months away. AI-assisted attacks have made the old threat timelines obsolete. This issue is about the new shape of enterprise security – and why it starts in the boardroom, not the SOC.

INTEL BRIEF

BY THE NUMBERS THIS WEEK

22 sec

Time for attackers to hand off from initial access to secondary threat actor – down from 8 hours 3 years ago

Google Cloud M-Trends 2026

28.3%

CVEs now exploited within 24 hours of public disclosure – before most patch cycles begin

Mandiant M-Trends 2026

1 in 3

Indian enterprises with no Zero Trust framework – despite 93% saying AI will strengthen their security

Zoho / Tigon Advisory Report, May 2026

₹ 250 Cr

Maximum penalty per contravention under India's DPDP Act Phase 2, effective November 2026

DPDP Phase 2 Framework, MeitY 2025



The Security Stack Is Being Rebuilt for the AI Era

Two major moves this week signal that legacy cybersecurity architecture is no longer fit for an agentic world

MAY 7, 2026

Cognizant Launches Secure AI Services — The Provable Trust Play

Cognizant launched Secure AI Services — spanning model security, AI DevOps security, agent behaviour controls, and generative AI risk management. The headline capability: Cognizant Neuro® Cybersecurity, a unified control plane for AI and enterprise signals. Already deployed across 250+ global enterprises. For CISOs: your next vendor conversation just changed.

GOOGLE CLOUD NEXT '26

Google + Wiz: The \$32B Cloud Security Reset

Google's \$32B acquisition of Wiz now reshapes cloud security architecture. At Cloud Next '26, Wiz announced support for AWS Agentcore, Gemini Enterprise, and Salesforce Agentforce — unified visibility across all AI platforms. The Triage and Investigation agent processed 5 million alerts last year, reducing 30-minute analysis to 60 seconds. For CTOs: the fragmented point security era is ending.

Sources: Cognizant Press Release, May 7 2026 · Google Cloud Blog — Next '26 Security

4×

Faster exfiltration speeds in 2025's fastest attacks vs prior year

Palo Alto Unit 42 2026

15 min

After a CVE is disclosed, attackers begin scanning — within 15 minutes

Palo Alto Unit 42 2026

90%

Identity weaknesses played a role in nearly 90% of Unit 42 incident investigations

Palo Alto Unit 42 2026

**RETINA
TAKE**

Your security stack was built for a world where attackers had days. They now have seconds. The enterprises winning the security battle are those that have moved from perimeter defence to identity-centric, AI-augmented continuous monitoring — with human analysts directing, not drowning in, the alert stream.



1

BOARD
LEVEL

The Attack Timeline Has Gone Negative — Your Patch Cycle Cannot Keep Up

Mandiant's M-Trends 2026 report delivers a finding every CTO and CISO must table at their next board meeting: 28.3% of CVEs are now being exploited within 24 hours of public disclosure — exploits routinely arrive before patches. The average time to remediate a known high-severity CVE is 74 days, while 45% of vulnerabilities in large enterprise systems never get remediated. The implication is structural: you cannot patch your way to security. For CTOs and CISOs: the priority shift is from remediation to resilience — continuous monitoring, threat containment, and assumption of breach as baseline posture.

Source: Mandiant M-Trends 2026 / Edgescan 2025 Vulnerability Statistics Report

2

INDIA
ALERT

DPDP Phase 2 Is Six Months Away — Most Indian Enterprises Are Not Ready

India's Digital Personal Data Protection Act Phase 2 becomes enforceable in November 2026 — activating Consent Managers, intensifying oversight of Significant Data Fiduciaries, and opening the 250 crore per-contravention penalty framework. The DPDP Act applies extraterritorially — any company offering goods or services to Indian individuals must comply, regardless of where it is based. AI systems processing Indian user data — voice AI, recommendation engines, analytics platforms — fall squarely within its scope. For GCs, CFOs, and CDOs: your DPDP readiness clock is not 18 months. It is 6 months.

Source: MeitY DPDP Rules 2025 / DLA Piper Data Protection Laws of the World 2026

3

RISK
GAP

India's Zero Trust Paradox: 93% Say Yes to AI Security, 1 in 3 Has No Zero Trust Plan

Zoho Corporation's State of Workforce Password Security Report 2026 exposes a critical contradiction: 93% of Indian businesses believe AI will strengthen their cybersecurity — yet 1 in 3 organisations have no Zero Trust strategy, and 34% report only partial visibility and control over critical business identities. Organisations are deploying AI-driven operations without full oversight of access privileges — exactly the unmonitored attack surface adversaries exploit. For CISOs and CHROs: Zero Trust is not an infrastructure project. It is an organisational discipline.

Source: Zoho / Tigon Advisory State of Workforce Password Security Report, May 2026



The DPDP Reckoning: Why India's Privacy Law Is a Board Problem, Not a Compliance Problem

By **Dr. Shinu Abhi** · Director – Corporate Training, REVA Academy for Corporate Excellence

Every few years, a regulatory shift arrives that organisations treat as a legal department problem — until it becomes a financial one. GDPR was that moment for Europe in 2018. India's DPDP Act is that moment for us in 2026.

Here is what most Indian enterprises have not yet internalised: the Digital Personal Data Protection Act does not care whether your privacy policy is beautifully written. It cares whether your systems can prove they honour it — in real time, with audit trails, across every AI model, voice platform, analytics pipeline, and third-party vendor that touches personal data. The Data Protection Board of India is now constituted. Phase 2 enforcement begins November 2026. The question is no longer "are we compliant?" but "can we demonstrate compliance on demand?"

Three consequences most boards have not yet mapped: First, AI changes the consent calculus entirely. A voice AI system that captures a customer's spoken response is processing biometric data under the Act — requiring explicit, purpose-specific consent. Second, third-party liability extends upstream. If your analytics vendor processes Indian user data outside India without an approved adequacy framework, you are the data fiduciary — and you are liable. Third, deletion is now operational, not theoretical. The right to erasure requires architecture that can locate, isolate, and delete personal data across distributed systems — including data used to train your AI models.

Sources: DPDP India 2026 Handbook (Levo.ai) · DPDP Phase 2 Compliance Guide (SecurePrivacy) · The Cyber Express: DPDP 2026 Impact

"2026 is the inflection year when planning becomes measurable operational work and when regulators will expect visible progress. It is no longer about intent — it is about evidence of execution."

Shashank Bajpai

CISO & CTSO, YOTTA · The Cyber Express, 2026

The enterprises that will navigate DPDP without disruption are those that have treated it as an engineering problem — building consent management, data inventory, access governance, and deletion workflows into their systems architecture, not their legal documentation. For everyone else, November 2026 will arrive as a shock.

RACE PERSPECTIVE

The DPDP Act is not a compliance cost. It is a forcing function for the data infrastructure discipline that Indian enterprises should have built three years ago. Organisations that execute well will emerge with stronger data governance, higher customer trust, and a competitive advantage in every regulated sector. The penalty is not the risk. The risk is being unable to operate with Indian user data after November 2026.

Deep Algorithms: The Indian Startup Building Identity for the Age of AI Agents

Case Study: Behavioural Identity at Canara Bank, Karnataka Bank, and Beyond

Source: *Indian Startup News — Deep Algorithms raises 16 Cr, April 2026* <https://deepalgorithms.io>

THE PROBLEM THEY SOLVE

AI agents act on behalf of users — but traditional IAM systems can't tell the difference between a human and an agent. Deep Algorithms builds identity systems that watch what an identity does, not just who it claims to be.

THE PRODUCT

adapID-AI — an adaptive behavioural intelligence engine that continuously models identity behaviour across systems, detecting credential misuse, privilege abuse, and AI agent anomalies in real time.

THE TRACTION

16 Cr raised from Unicorn India Ventures. 3× ARR growth from FY25 to FY26. Clients: Canara Bank, Karnataka Bank, DCB Bank, CSB, NCRTC. Expects 20× ARR by FY2027.

THE FRONTIER MOVE

Launching adapID AI CIRM — the first purpose-built Agentic Identity Risk Management platform. The product category did not exist 18 months ago.

What This Means for CXOs

CISO

Your IAM vendor built their product for human logins. Every AI agent is an unmonitored identity with potentially over-privileged access. Deep Algorithms addresses exactly this gap.

CTO

Behavioural identity is the next frontier of access control. Static role-based permissions will not survive agentic workflows. Evaluate vendors who monitor machine identities dynamically.

CFO

BFSI clients like Canara Bank and Karnataka Bank are deploying this because the cost of an identity breach in regulated banking far exceeds the cost of prevention.

CEO

India is producing enterprise-grade AI security IP. A Hyderabad startup is defining a product category that global incumbents are scrambling to catch up with. Watch this space.

**Sridhar Govardhan**

Group Chief Information Security Officer Group, Angel One
Mentor, Cybersecurity, RACE

Q1

The DPDP Act enforcement is six months away. What is the single most dangerous assumption Indian enterprises are making right now?

A

The most dangerous assumption is that readiness equals documentation. It's easy to believe that having a privacy policy, a data map, and a DPO in place means you're prepared, and honestly, that's where most organisations start. But what the DPDP Act demands is operationalised compliance, live consent mechanisms, grievance redressal that actually functions, and data processors who are contractually and technically accountable. The gap between documented intent and operational reality is where the real work lies. The question every enterprise should be asking isn't "are we compliant on paper?" It's "have we stress-tested our compliance in practice?"

The Board needs to ask one question: "If our regulator asked us to prove our DPDP compliance tomorrow morning, what would we show them?" If that question creates silence in the room, you are six months from a serious problem.

Q2

You've spent two decades in Cybersecurity. How has AI changed what a breach actually looks like?

A

In two decades, the most significant shift AI has brought is the collapse of the dwell-time advantage defenders once had. Breaches used to follow recognisable patterns. We have time to detect, respond, and contain. AI-powered attacks compress that window dramatically. We're also seeing a blurring of the perimeter: with AI-generated deepfakes, voice cloning, and hyper-personalised social engineering, the human layer is now the most vulnerable surface in any enterprise. A breach today often doesn't start with a vulnerability; it starts with a conversation.

Q3

What is your one non-negotiable recommendation to a CISO preparing for an AI-first threat environment?

A

Invest in AI literacy for your security team before investing in AI security tools. The market is full of powerful AI-powered solutions, and many of them are genuinely valuable, but tools are only as effective as the people interpreting their outputs. A CISO preparing for an AI-first threat environment needs analysts who understand how large language models can be manipulated, how autonomous agents behave under adversarial conditions, and where AI introduces blind spots in detection logic. Technology is the easy part. Understanding it deeply enough to defend against its misuse, that's the non-negotiable.

Note: CANDID represents a representative editorial perspective compiled from public commentary. RACE welcomes direct contributions from CXO leaders. Contact: retina@reva.edu.in



5 Roles With Exploding Demand · Week of 11 May 2026 · Focus: Security & Data Governance

#	Role & Critical Skills	CTC Range	Source
1	AI Security Engineer LLM security · Prompt injection defence · AI red teaming · DevSecOps	₹30L–₹75L	Taggd India AI Engineer Salary Report 2026*
2	Data Protection Officer (DPO) DPDP Act · GDPR · Data mapping · Privacy by Design · Risk assessment	₹35L–₹80L	upGrad Cybersecurity Salary Report 2026**
3	IAM Architect Zero Trust · PAM · Machine identity · SailPoint · CyberArk	₹28L–₹65L	Networkers Home Cloud Security Salary India 2026***
4	SOC Analyst (AI-Augmented) SIEM/SOAR · AI alert triage · Threat hunting · Incident response	₹12L–₹30L	India Decoding Jobs 2026 (Taggd × CII)#
5	GRC Lead (AI Governance) EU AI Act · ISO 42001 · Risk frameworks · Compliance automation	₹22L–₹50L	upGrad Cybersecurity Salary Report 2026##

Sources:

*AI Security Engineer & DPO: Taggd India AI Engineer Salary Report 2026

**upGrad Cybersecurity Salary Report 2026.

***IAM Architect & SOC Analyst: Networkers Home Cloud Security Salary India 2026

#India Decoding Jobs 2026 (Taggd × CII).

##All ranges reflect mid-to-senior professional profiles in major Indian metros.



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

This week's Research Radar features four new RACE publications spanning cloud resilience, AI engineering, intelligent sensing, and software quality. None repeat from Issue #1. All are accessible via the links below.

CLOUD · AGENTIC AI · AIOPS

Autonomous Multi-Agent System for Integrated SRE and Self-Healing in Cloud-Native Environments

N. V. Manohar, Ravi Shukla & Dr. Shinu Abhi ·
ICSSSM 2025, Atlantis Press / Springer Nature ·
Dec 2025

CrewAI + LLM (DeepSeek-R1 / GPT-4o) framework automating the full incident management lifecycle. Results: 96% RCA accuracy, 73% automated code-fix success rate, MTTR reduced from hours to under 28 minutes.

atlantispress.com/proceedings/icsssm-25/126020167

AI ENGINEERING · SOFTWARE QUALITY

LLM Driven Unit Test Case Generation Using Agentic AI

Baskaran S., Pradeepta Mishra & Rashmi Agarwal · Journal of Ubiquitous Computing & Communication Technologies, Vol. 7 No. 4 · Nov 2025

Autogen Agentic AI + LLMs to auto-generate, validate, and improve Python unit tests. Mutation score improved from 83.9% to 95.8% on an Insurance Management Application.

irojournals.com/jucct/article/view/1818

IOT · HEALTHCARE · DEEP LEARNING

Fall Detection Using Wi-Fi Signals: Adaptive Residual Dense Network with CBAM and Bio-Inspired Optimisation

Dheeraj Sharma, J. B. Simha & Rashmi Agarwal · Springer · Jan 2026

Non-intrusive fall detection using Wi-Fi CSI signals converted to 2D images. ENBOA-ARD-CBAM achieves F1 score 10% better than ResNet and 5% better than VGG-16. ROC exceeds 0.8

link.springer.com/chapter/10.1007/978-981-96-7545-6_3

HEALTHCARE AI

Medical Bad Debt Propensity Prediction: Comparative Analysis using ML Algorithms

Anushree BV, Yuvaraju M & Dr. Shinu Abhi · ICMSCI 2026, IEEE Xplore

Trained on 182,980 hospital encounters from MEDITECH EHR (2021–2025). LightGBM achieved AUC-ROC 0.884 and concentrates 51.4% of bad debt cases in the top risk decile — enabling revenue cycle intervention at discharge, not at 120–180 day collection triggers. XGBoost captured 93.6% bad debt recall. Targets 15–20% AR recovery improvement.

ieeexplore.ieee.org/document/11469311


**RETINA
TAKE**

RACE research is moving from academia into direct enterprise application — self-healing cloud systems, AI-augmented software quality, and IoT-based health safety. Bad Debt Predictions are not theoretical exercises. They are live deployment problems these papers are solving. Explore the full RACE research archive at racelit.netlify.app



BOOK SHELF

RETINA'S READ OF THE WEEK



RETINA RECOMMENDS · THIS WEEK'S THEME

The CISO Evolution: Business Knowledge for Cybersecurity Executives*Matthew K. Sharp, Kyriakos P. Lambros · Wiley, 2022*

As the CISO role transforms from technical gatekeeper to strategic board voice, the competency gap has never been wider. The authors map exactly what a modern security leader needs beyond technical mastery — financial literacy, board communication, business risk framing, and the language of the C-suite. In a week when India's enterprises are being asked to treat cybersecurity as a boardroom discipline rather than an IT function, this is required reading for any CISO preparing their next board presentation — and for every CEO evaluating whether their CISO can make the shift.



RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

SATURDAY, 23 MAY 2026 · 9:30 AM IST

Namma Bengaluru Meetup Mastering Agentic AI

From Intelligent Agents to Secure Deployment. Practitioner-led session on building, governing, and securing multi-agent systems in enterprise environments. Open to all.

Contact race@reva.edu
Saturday, 23 May 2026
 09:30 AM Onwards IST
REGISTER FREE

SATURDAY, 16 MAY 2026 · 4:30 PM IST

REVA Corporate Cricket League 2026

The second edition of RCL brings together RACE alumni, corporate partners, and enterprise teams for an afternoon of cricket and connection at REVA Cricket Ground, Bengaluru.

**SCAN ME**

Live non-stop cricket score
on #CricHeroes.

UPCOMING · SEPTEMBER 2026 BATCH

Applications are now open for all six RACE Master's & PG Programmes — Business Analytics (#1 India), Artificial Intelligence (#3 India), Cybersecurity (#5 India), and Cloud Architecture & Security. Weekend format. AICTE/UGC approved. Industry mentors from Amazon, Microsoft, AWS, EC-Council, Terralogic, IBM and more.

Apply before 31 July 2026. race.reva.edu.in

RACE PROGRAMMES

REVA Academy for Corporate Excellence · REVA University · Bengaluru · AICTE / UGC Approved · Weekend Batches · Industry Mentors

Master's & PG Programmes

New Batch: September 2026 · 12–24 months · Weekend · For Working Professionals

#1 IN INDIA · ANALYTICS INDIA MAGAZINE ANALYTICS PG Diploma / M.Sc. in BUSINESS ANALYTICS AWS Academy · Microsoft Azure 12–24 Months	#3 IN INDIA · TECHGIG ARTIFICIAL INTELLIGENCE M.Tech. in ARTIFICIAL INTELLIGENCE AWS Academy · Microsoft Azure 24 Months	#3 IN INDIA · TECHGIG ARTIFICIAL INTELLIGENCE PG Diploma / M.Sc. in ARTIFICIAL INTELLIGENCE AWS Academy · Microsoft Azure 12–24 Months
#5 IN INDIA · ANALYTICS INDIA MAGAZINE CYBERSECURITY M.Tech. in CYBERSECURITY EC-Council · Microsoft Azure · Terralogic Inc. 24 Months	#5 IN INDIA · ANALYTICS INDIA MAGAZINE CYBERSECURITY PG Diploma / M.Sc. in CYBERSECURITY EC-Council · Microsoft Azure · Terralogic Inc. 12–24 Months	CLOUD PG Diploma / M.Sc. in CLOUD ARCHITECTURE & SECURITY AWS Academy · Microsoft Azure 12–24 Months

Short-Term Certification Programmes

100 Hours · Project-Driven · Customisable for Organisations

LEADERSHIP 8–12 weeks TECHNO-FUNCTIONAL LEADERSHIP Executive presence, decision-making, leading tech-driven teams. Built for mid-to-senior managers.	CYBERSECURITY 8–12 weeks CYBERSECURITY & SOC OPERATIONS Threat detection, incident response, SIEM/SOAR. Hands-on in RACE's live Security Operations Centre.
AI ENGINEER 8–12 weeks AI ENGINEERING CERTIFICATION Data pipelines, model training, LLMs, deployment. Real production projects every batch.	CLOUD 8–12 weeks CLOUD ARCHITECTURE & DEVSECOPS Multi-cloud, IaC, secure CI/CD on real AWS/Azure/GCP environments. Not demos.
AGENTIC AI 4–6 weeks AGENTIC AI SYSTEMS DESIGN Multi-agent systems using LangChain, LangGraph, CrewAI. Capstone: deployed autonomous agent.	DATA & AI 12 weeks DATA SCIENCE & BUSINESS ANALYTICS Statistics to ML to BI. Python, SQL, Power BI, Tableau on real business datasets.

Enroll with us · race@reva.edu · race.reva.edu.in

This Week, Ask Your Team

"Can you give me a complete inventory of every AI agent currently
operating in our organisation — **who owns it, what data it can
access, and what happens if it fails?**"

RETINA PUBLISHES EVERY SUNDAY · FORWARD TO A CXO YOU RESPECT