

RACE WEEKLY INTELLIGENCE BRIEF

RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER

The Agentic Enterprise: AI Takes the Helm

ALSO INSIDE

Sarvam AI · Post-Quantum · CISO Elevated

FROM THE EDITOR'S DESK

Welcome to Issue #1 of RETINA – RACE's weekly intelligence brief crafted exclusively for India's enterprise leaders. Each Sunday, we distil the week's most consequential technology signal into a focused 10-minute read. This week: AI agents stopped being a promise and became infrastructure. We tell you what that means for your organisation, your boardroom, and your next hire.

HYPERFOCAL

INNOVATIONS & INVENTIONS
WHAT BROKE THIS WEEK

The Agent Is In: AI Takes the Enterprise Helm

Two platform announcements confirm agentic AI has crossed from experiment to enterprise standard

Two announcements this week confirmed what analysts have been predicting since 2024: AI agents have crossed the chasm from experiment to enterprise standard. The implications for every CXO in the room are immediate.

Adobe CX Enterprise – The Agentic Marketer

Adobe retired its 12-year-old Experience Cloud brand this week, replacing it with CX Enterprise – an AI-first platform where 'Coworker' agents orchestrate the entire customer lifecycle autonomously. Persistent agents now handle brand compliance, content production, journey personalisation, and loyalty delivery without human prompting on each step. For CMOs and CDOs: this is not an upgrade. It is a rearchitcting of how marketing operates.

OpenAI GPT-5.5 & Workspace Agents – The Super App Vision

OpenAI launched GPT-5.5 and simultaneously introduced Workspace Agents – AI units shared across teams that autonomously gather context, execute multi-step tasks in Slack, Gmail, and Drive, and seek approval when needed. The model positions ChatGPT less as a chat interface and more as an operating system for knowledge work. For CTOs and COOs: the question is no longer 'should we try AI agents?' but 'do we have governance for the agents that are already being deployed?'

84% Reduction in case resolution time reported by Salesforce Agentforce at Reddit Salesforce Internal Data, 2026	\$100M Annual operational savings from Agentforce – the benchmark every COO will now face Salesforce Internal Data, 2026	13% Organisations that have successfully integrated AI across the full SDLC despite 90% using AI JetBrains AI Pulse Survey, Jan 2026
---	---	---

RETINA
TAKE

The race is no longer about model capability – it is about enterprise orchestration, governance, and integration. The CXOs who will win are those who define agent ownership, accountability structures, and human escalation protocols now, before their boards ask why they didn't.



1

BOARD
LEVEL

CISO Elevation: Security Has a Boardroom Seat — Permanently

Gartner research confirms cybersecurity leadership is undergoing structural transformation. The 'unicorn CISO' who handles both governance and engineering is disappearing. Forward-thinking enterprises are bifurcating the role into a strategic CISO (enterprise risk, board relations, regulatory compliance) and a VP of Security Engineering (technical execution). Boards now treat cyber incidents with the same financial scrutiny as liquidity ratios. For CEOs and CHROs: this is a hiring signal. The next CISO you appoint is a business executive, not a technology one.

Source: Gartner Top Cybersecurity Trends 2026

2

URGENT

Post-Quantum Cryptography Moves from 'Someday' to 'This Year'

Gartner predicts advances in quantum computing will render current asymmetric encryption standards unsafe by 2030. With 'harvest now, decrypt later' attacks already targeting long-lived sensitive data, the window for enterprises to achieve cryptographic agility is narrowing rapidly. 94% of executives in the WEF Global Cybersecurity Outlook 2026 identified AI as the most significant driver of cyber risk change, while only 64% have processes to assess AI security — up from 37% last year. For CTOs and CISOs: post-quantum migration planning is no longer optional. Start your cryptography inventory now.

Sources: Gartner Feb 2026 · WEF Cybersecurity Outlook 2026

3

RISK

AI Governance Gaps Are Becoming Regulatory Liability

The EU AI Act is operational. US state-level AI laws are proliferating. Yet most enterprises have deployed AI tools without formal governance frameworks. Gartner's analysis shows that as AI agents proliferate, IAM systems built for human users are dangerously inadequate for machine actors — creating unauthorised access risks at scale. For GCs and CFOs: the cost of non-compliance is no longer theoretical. The question to ask your CTO this week: *'Can you give me a full inventory of every AI agent operating in our enterprise?'*

Source: Gartner Agentic AI & IAM Report 2026



The Governance Gap: Why 2026 Is the Year Enterprises Must Own Their AI Agents

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

Here is the uncomfortable truth sitting in most enterprise boardrooms in April 2026: your organisation almost certainly has AI agents operating in your environment that no one has formally approved, governed, or audited.

April 2026 marks a structural inflection point in enterprise AI adoption. Agentic orchestration — where AI systems execute complex, multi-step workflows autonomously — has moved from pilot to production at organisations like EY (processing 1.4 trillion lines of audit data annually through its Canvas platform), Salesforce, and JPMorgan. These are not technology experiments. They are mission-critical infrastructure.

But here is the governance gap that should concern every CXO: while the tooling for agent deployment has become democratised — no-code platforms, low-code agent builders, GPT-based workspace agents — the governance infrastructure has not kept pace. Traditional Identity and Access Management (IAM) systems were built for human identities. AI agents require a fundamentally different identity and policy architecture.

“The best engineering organisations of the next decade will not be defined by how many engineers they have, but by how effectively human judgment and AI capability work as one.”

Rajesh Varrier

President – Operations & Chairman, MD, Cognizant India · April 2026

The enterprises that will thrive are those that move on three fronts simultaneously. First, complete an AI agent inventory. Second, extend your IAM architecture to include machine actors with the same rigour applied to human identities. Third, build incident response playbooks specifically for agent failures, not just human errors.

RACE PERSPECTIVE

The organisations that will struggle in 2027 are those treating AI governance as a compliance checkbox. The organisations that will lead are those treating it as competitive infrastructure — because the speed at which you can safely deploy, govern, and iterate on AI agents will determine your operational velocity for the next decade. why they didn't.

Sarvam AI: Building India's Sovereign Intelligence

Case Study: From IIT Lab to National AI Infrastructure

THE FOUNDERS

Vivek Raghavan and Pratyush Kumar — creators of AI4Bharat at IIT Madras — left academia to build full-stack AI natively for India's 1.4 billion people across 22 languages.

THE MODEL

A sovereign AI stack: small language models (Sarvam-2B), voice AI (Bulbul V3 with 35+ voices in 11 languages), ASR for 22 Indian languages, and Sarvam Agents for enterprise workflows.

THE TRACTION

\$53M Series A from Lightspeed, Peak XV, and Khosla Ventures. Selected by the Government of India under the IndiaAI Mission to build India's first sovereign LLM. Unveiled Sarvam-30B and Sarvam-105B at the India AI Impact Summit 2026.

THE EDGE

Vision OCR scored 84.3% on olmOCR-Bench — beating Gemini 3 Pro (80.2%) and ChatGPT (69.8%) — on multilingual document understanding. Built for Bharat-scale.

What This Means for CXOs

CEO

India-first AI is no longer a compromise. Sovereign models that understand your customers in their own language may outperform global platforms on your most critical use cases.

CTO

Evaluate small language models (SLMs) alongside large ones. Sarvam-2B running on-premise in a regulated environment may be more viable than GPT-5 via API.

CHRO

Voice AI in regional languages is unlocking Tier 2 and Tier 3 talent pools. Bajaj Finance now disburses \$600M in loans annually almost entirely through voice AI — powered by an Indian startup.

CFO

Sovereign AI reduces data residency risk and regulatory exposure. The compliance cost of sending sensitive data to foreign models is rising — factor it into your AI TCO.



Sandeep Vijayaraghavan

EVP, Cybersecurity & Cloud, Terralogic
Chief Mentor, Cybersecurity, RACE

Q1

The word 'agentic AI' is everywhere this week. For a CXO who isn't from a technology background, what does it actually mean for day-to-day business?

A

For a CISO, this means AI systems are no longer passively responding to prompts; they're autonomously executing multi-step tasks — browsing, coding, emailing, and querying systems — with minimal human oversight. Adoption is outpacing most security teams' ability to track, with agentic AI quietly embedding itself into IT ops, finance, and software pipelines almost overnight.

Your security controls were built for humans. Non-human identities — AI agents acting on behalf of users — don't fit neatly into your IAM policies, don't trigger your usual access alerts, and often inherit far more permissions than they should. The attack surface just got a lot harder to see.

Q2

What is the single biggest mistake Indian enterprises are making right now with AI adoption?

A

In the rush to show AI ROI to the board, security teams are being brought in after the business has already committed — leaving CISOs inheriting risk they had no hand in designing. In a market where data localisation, regulatory scrutiny from DPDP, and third-party vendor risk are all converging at once, that sequencing mistake isn't just a technical gap — it's a boardroom liability waiting to happen.

Q3

What would you tell a CHRO who is anxious about AI and workforce impact?

A

The threat isn't AI — it's standing still. The real risk isn't that AI replaces your security team — it's the myth that AI is infallible. AI will absorb the repetitive — the alerts, the reports, the pattern-matching. What it cannot absorb is your team's judgment, instinct, and accountability. The security professional who understands AI's limits — and works with it critically — will always outpace one who either fears it or blindly trusts it. Fluency beats fear. Every time

5 Roles With Exploding Demand · Week of 27 April 2026

#	Role & Critical Skills	CTC Range	Source
1	GenAI / LLM Engineer LangChain · RAG · LLM fine-tuning · Python · MLOps	20L– 70L	Taggd India AI Engineer Salary Report 2026
2	CISO (Strategic) Risk governance · ISO 27001 · Zero Trust · Board comms	50L– 1Cr+	upGrad Cybersecurity Salary Report 2026
3	Cloud Security Architect AWS/Azure Security · Zero Trust · DevSecOps	30L– 60L	Networkers Home Cloud Security Salary India 2026
4	AI / ML Engineer (Senior) Python · TensorFlow · MLOps · GenAI deployment	25L– 60L	ixceed-solutions India 2026 Salary Data
5	DevSecOps Engineer Kubernetes security · CI/CD pipelines · Container scanning	18L– 40L	Networkers Home Cloud Security Salary India 2026

Sources: Salary ranges for mid-to-senior professionals. GenAI/LLM Engineer & AI/ML Engineer: Taggd India AI Engineer Salary Report 2026. CISO: upGrad Cybersecurity Salary Report 2026. Cloud Security Architect & DevSecOps Engineer: Networkers Home Cloud Security Salary India 2026. Demand context: India Decoding Jobs 2026 (Taggd × CII).



This week, RACE researchers published five applied AI papers at the Second International Conference on Multi-Agent Systems for Collaborative Intelligence (ICMSCI), with IEEE Xplore links now live. The work spans enterprise AI, cybersecurity, healthcare analytics, demand forecasting, and computer vision for railway safety.

ENTERPRISE AI

Bridging Enterprise Silos: Multi-Agent Bot for Sales Enablement

Generative AI, CrewAI and RAG for improving sales enablement and decision support.

ieeexplore.ieee.org/document/11469345

CYBERSECURITY

Detection and Reporting of Typosquatting Malicious Domains with Machine Learning

XGBoost-based system for detecting phishing and malicious domains.

ieeexplore.ieee.org/document/11469751

HEALTHCARE ANALYTICS

Medical Bad Debt Propensity Prediction: Comparative Analysis using ML Algorithms

LightGBM-based prediction model for early hospital revenue-risk intervention.

ieeexplore.ieee.org/document/11469311

DEMAND FORECASTING

Machine Learning and Hybrid Approaches for Intermittent Demand Forecasting

Random Forest-based forecasting for intermittent spare-parts demand.

DOI: 10.1109/ICMSCI67830.2026.11469346

COMPUTER VISION · SAFETY

Multi-Task Safety Monitoring for Railway Crossings using Deep Learning and Computer Vision

YOLOv8-based vision system for gate-status classification and obstacle detection.

ieeexplore.ieee.org/document/11469541

RETINA TAKE

RACE's publication pipeline is moving strongly toward AI systems that solve real enterprise and societal problems — from sales productivity and cyber defence to healthcare finance and transport safety.

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

NAMMA BENGALURU · MEETUP

MASTERING AGENTIC AI
From Intelligent Agents to Secure Deployment.



 **Saturday, 23 May 2026**
09:30 AM Onwards IST

REGISTER FREE


RCL 2026
REVA CORPORATE CRICKET LEAGUE
Saturday, 16 May 2026 | 4:30 PM onwards
VENUE: REVA CRICKET GROUND




SCAN ME
Live non-stop cricket score on #CricHeroes.