

RACE WEEKLY INTELLIGENCE BRIEF

RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER:

The ROI Reality Check: The Pilots Are Over. The Bill Has Arrived.

ALSO INSIDE

Yotta's Sovereign Compute Bet · REGWATCH

Three Deadlines, One Quarter · The Board Pack · The 90-Day Value Playbook

FROM THE EDITOR'S DESK

Issue #1 asked who owns your AI agents. Issue #2 asked who secures them. Issue #3 asked who owns the data beneath them. Issue #4 asked whether your leaders understand what they approve. Issue #5 asked what AI is doing to your people. This week, we arrive at the question your CFO has been waiting to ask: where is the return? Three years of enterprise AI investment are now entering their first serious budget review cycles – and the evidence says most organisations cannot answer that question with a number. This issue is about why, and about the measurement discipline that separates the enterprises compounding value from those quietly cancelling programmes.



SIGNAL

SIGNAL THE LITERACY AGENDA – BY THE NUMBERS

95%

of generative AI pilots delivered zero measurable P&L impact - value realisation, not model quality, is the failure point

MIT Project NANDA, State of AI in Business

42%

of companies abandoned most of their AI projects in 2025 - the first wave of budget reviews is cancelling what cannot be measured

S&P Global, 2025–26

21%

of S&P 500 companies could cite any measurable benefit from their AI investments

Morgan Stanley Research, 2026

\$19.8 Bn

JPMorgan's 2026 technology budget – with AI formally reclassified from experimental R&D to core infrastructure

JPMorgan Chase, June 2026



JPMorgan Moves AI From 'Experiment' to 'Core Infrastructure' – and Resets the Benchmark for Every Board

JPMorgan Chase formally reclassified its AI investment from experimental R&D to core infrastructure, anchored by a 2026 technology budget of roughly \$19.8 billion and around 2,000 staff dedicated to AI development. Meta, in parallel, signalled AI capital expenditure of \$115–135 billion for 2026, nearly double the prior year. The signal for Indian enterprise boards is not the spending itself. It is the accounting posture: when the world's most scrutinised bank stops treating AI as an experiment, 'pilot mode' stops being a defensible position in any boardroom. For CFOs: if AI sits in your innovation budget rather than your operating plan, your governance, measurement, and accountability structures are one cycle behind your peers.

Source: <https://www.aicerts.ai/news/jpmorgan-makes-ai-part-of-core-bank-infrastructure/>

The Enforcement Quarter Begins: Colorado on June 30, the EU on August 2

The regulatory honeymoon ends this month. The Colorado Consumer Protections for Artificial Intelligence Act takes effect on June 30, 2026 – the first substantive AI enforcement deadline on US soil, covering high-risk AI in employment, healthcare, financial services, education, housing, and legal services. It requires risk management programmes, annual impact assessments, disclosure when AI drives consequential decisions, and a right of appeal for affected residents. Five weeks later, on August 2, the EU AI Act reaches full application, with high-risk system rules and transparency obligations enforceable. Indian enterprises serving US or EU customers – every GCC, every IT services firm, every SaaS exporter – are in scope. For GCs and CTOs: ROI now has a compliance denominator. An AI system that cannot pass an impact assessment is not an asset; it is a liability with a go-live date.

Source: Colorado AI Act (effective 30 Jun 2026): <https://leg.colorado.gov/bills/sb24-205>

RETINA TAKE

The era of AI spending without AI accounting is over. Between JPMorgan's reclassification and the June–August enforcement window, boards now face the same question from two directions - regulators asking 'can you prove control?' and CFOs asking 'can you prove return?' The organisations that built measurement infrastructure in 2025 can answer both. Everyone else is about to discover the cost of not being able to.

 PANORAMA

THREE FORCES SHAPING YOUR DECISIONS

1

MEAS-
URE-
MENT**The 95% Problem Is a Measurement Failure, Not an AI Failure**

The most repeated statistic in enterprise AI – MIT's finding that 95% of pilots show no P&L impact – is being widely misread as proof that AI does not work. Researchers at UC Berkeley's Sutardja Center offer the sharper reading: most organisations apply traditional software payback logic to a technology that behaves nothing like traditional software, and they track usage metrics – seats, logins, hours – that say nothing about outcomes. Roughly 80% of the work of moving from pilot to production is data engineering, governance, workflow integration, and measurement infrastructure; most pilots launch with none of it and no predefined success criteria, which means there is no way to declare success even when the technology performs exactly as designed. For CEOs: before asking whether your AI works, ask whether you could detect it if it did.

Source:

- MIT Project NANDA – State of AI in Business (95% zero P&L finding)
https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf
- UC Berkeley Sutardja Center – measurement-failure reframe <https://scet.berkeley.edu/>

2

BUDGET
CYCLE**The Bill Comes Due: 2022–23 Pilots Are Hitting Their First Real Budget Reviews**

Organisations that launched AI pilots in 2022 and 2023 are now inside their first serious budget review cycles – and programmes that cannot demonstrate measurable return are being cancelled at scale. S&P Global found 42% of companies abandoned most of their AI projects in 2025; IBM puts initiatives delivering expected ROI at roughly one in four. KPMG's Global AI Pulse Survey shows the gap is widening rather than closing: among the small group of AI leaders who moved beyond pilots to scaled deployment, 82% report AI already delivering meaningful business value, against 62% of peers – and three in four leaders say they will prioritise AI investment even through economic uncertainty. The lesson is not to spend less. It is that capital is consolidating around the organisations that can prove value, and away from those that cannot. For CFOs: fund measurement before models. A smaller portfolio you can evidence beats a larger one you cannot.

Source: S&P Global – AI project abandonment data

<https://www.spglobal.com/market-intelligence/en/news-insights/research/ai-experiences-rapid-adoption-but-with-mixed-outcomes-highlights-from-vote-ai-machine-learning>

3

INDIA
EDGE**India Leads on Adoption. The ROI Race Will Be Won on Discipline.**

Deloitte's State of AI in the Enterprise 2026 confirms India's lead on adoption momentum: 40% of Indian respondents report significant or full AI usage against a global average of about 28%, with at-scale deployment strongest in product development, operations, marketing, and supply chain. But the same research flags that Indian organisations rank below global counterparts on high-level AI expertise – and Issue #5's workforce findings showed capability building lagging deployment. Put those together with this week's global ROI data and the strategic risk is clear: adoption leadership without measurement discipline converts into the world's largest portfolio of unevidenced AI spend. For Indian CEOs: the next phase of the India AI story will not be written in deployment percentages. It will be written in audited business outcomes.

Source: Deloitte State of AI in the Enterprise 2026 – India insights

<https://www.deloitte.com/in/en/about/press-room/indian-enterprises-lead-global-peers-in-at-scale-ai-adoption-across-most-functions.html>



REGWATCH

NEW · THE COMPLIANCE COUNTDOWN

A new fixture in RETINA: every dated AI and data obligation on the horizon, in one table. Hand this page to your GC.

DEADLINE	REGULATION	WHO IS EXPOSED	DO BY THEN
30 Jun 2026	Colorado AI Act (USA)	Any enterprise whose high-risk AI touches Colorado residents — GCCs, IT services, SaaS exporters	Risk management programme, impact assessments, consumer disclosure and appeal mechanism
2 Aug 2026	EU AI Act — full application	Providers and deployers of high-risk AI serving EU markets; extraterritorial reach	System classification, conformity assessment, transparency documentation, Article 4 literacy evidence
Nov 2026	India DPDP Act — Phase 2 enforcement	Every data fiduciary processing Indian personal data; 250 Cr per-contravention exposure	Live consent management, deletion architecture, processor accountability, breach readiness

Sources:

1. Colorado AI Act (effective 30 Jun 2026): <https://leg.colorado.gov/bills/sb24-205>
2. EU AI Act — full application 2 Aug 2026: <https://artificialintelligenceact.eu/implementation-timeline/>
3. India DPDP Act & Rules (MeitY): <https://www.meity.gov.in/data-protection-framework>



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

The ROI Question Indian Boards Are Asking Wrong — and the One They Should Ask Instead

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

There is a moment in nearly every executive cohort at RACE when the conversation turns to AI ROI, and a senior leader asks some version of the same question: 'How do we calculate the return on our AI investment?' It sounds like the right question. It is usually the wrong one - because it assumes the organisation has already done the thing that makes the calculation possible: defined, before deployment, what the AI was

supposed to change, by how much, measured against what baseline.

Most have not. In three years of working with enterprise cohorts, the pattern is remarkably consistent. AI initiatives are approved on capability demonstrations; what the system can do, rather than outcome commitments, what the business will measurably stop doing, do faster, or do better. Success criteria, where they exist at all, are usage metrics:



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

licences activated, prompts issued, hours logged. These numbers are easy to collect and comfortable to present. They are also irrelevant to the only question a board is entitled to ask: did this produce a better outcome than what it replaced?

The fix is not a more sophisticated ROI formula. It is sequencing. The enterprises that can evidence AI value in 2026 made three unglamorous moves before scaling anything. They tied every initiative to a specific, named business decision or process with a measurable baseline. They built the instrumentation — error rates, cycle times, rework, adoption depth — into the workflow itself, not into a quarterly slide. And they gave

a single accountable owner the authority to kill initiatives that miss their numbers, which is the discipline that makes the surviving numbers believable.

There is one more shift Indian boards specifically need to make. Our enterprises lead the world in AI adoption breadth, and that breadth is becoming a liability when it is unmeasured. Forty initiatives with no baselines will lose, every time, to eight initiatives with audited outcomes. The question for your next board review is not 'how much have we deployed?' It is: 'show me the five AI initiatives where we can prove the delta — and the plan to retire the rest or instrument them properly.'

RACE PERSPECTIVE

AI ROI is not a finance exercise performed after deployment. It is a design constraint imposed before it. The enterprises that will compound value through 2027 are those that treat every AI initiative as a falsifiable business hypothesis — with a baseline, an owner, an instrument, and a kill criterion.



METRIC OF THE WEEK

NEW · ONE NUMBER,
PROPERLY UNDERSTOOD

Time-to-Value (TTV)

What it is: the elapsed time from initiative approval to the first measured, baselined business outcome - not to deployment, not to go-live, to evidenced value.

Why it matters now: focused, workflow-level AI implementations are demonstrating value within 6–12 months; enterprise-wide programmes typically take 2–4 years. If an initiative's projected TTV exceeds your board's patience, that mismatch — not the technology — is what will kill it. Surface it at approval, not at review.

How leaders use it: portfolio balance. Pair long-TTV foundation bets (data, governance, platforms) with short-TTV workflow wins that fund organisational patience. A portfolio made only of long bets loses its budget; only of quick wins, its future.

 ZOOM

STARTUP CASE STUDY – THE INNOVATION LENS

Yotta's Shakti Cloud: The Economics of Sovereign Compute – and the TCO Question Every CFO Should Re-Run

Case Study: From Hyperscale Data Centres to India's Largest Sovereign GPU Cloud

THE COMPANY: Yotta Data Services, backed by the Hiranandani Group, has built Shakti Cloud – sovereign AI infrastructure powered by more than 20,000 NVIDIA Blackwell Ultra GPUs across hyperscale campuses in Navi Mumbai and Greater Noida, delivered on a pay-per-use model and aligned with the IndiaAI Mission.

THE THESIS: Make industrial-scale AI training and inference affordable and compliant for Indian enterprises and the public sector – GPU-dense, high-bandwidth compute inside Indian jurisdiction, priced for consumption rather than commitment.

THE ROI ANGLE: Compute is the largest variable cost line in most enterprise AI P&Ls – and the line most often estimated once and never revisited. Domestic GPU capacity at this scale changes inference economics, data residency cost, and DPDP compliance overhead simultaneously. AI TCO models built on 2024 assumptions about foreign cloud pricing are stale.

What This Means for CXOs

CFO: Re-run your AI TCO with current domestic compute pricing before approving the next platform commitment. The sovereign premium you priced in two years ago may now be a discount.

CTO: Pay-per-use sovereign GPU capacity makes workload portability a negotiating lever. Architect for it; do not let inference lock-in quietly accumulate.

CHRO: Domestic data centre buildout at multi-gigawatt scale is creating a new talent market in AI infrastructure operations – a hiring pool and an attrition risk at once.

BOARD: Sovereign compute converts a compliance cost into a strategic option. Boards that asked 'cloud vs on-prem' in 2024 should be asking 'global vs sovereign vs hybrid' in 2026 – with numbers.

Sources:

1. NVIDIA India AI Mission briefing <https://blogs.nvidia.com/blog/india-ai-mission-infrastructure-models/>
2. Shakti Cloud <https://shakticloud.ai>



THE BOARD PACK

NEW · TABLE THIS AT YOUR NEXT REVIEW

- **ASK No. 1:** Which of our AI initiatives have a pre-deployment baseline and a named accountable owner - and what happens to the ones that do not?
- **ASK No. 2:** If our regulator - or our largest customer - asked us tomorrow to evidence the controls and the business impact of our highest-risk AI system, what would we show them?
- **ASK No. 3:** What did we retire this quarter? An AI portfolio where nothing is ever killed is a portfolio nobody is measuring.

TRACK

Evidenced-Value Ratio - the share of your live AI initiatives that can demonstrate a measured improvement against a pre-deployment baseline. Most organisations discover theirs is below 20%. The leaders run above 60% - not because more projects succeed, but because fewer unmeasured ones survive.



THE 90-DAY PLAYBOOK

NEW · FROM UNMEASURED
TO EVIDENCED

A practical sequence to stand up AI value measurement in one quarter — built from the patterns RACE observes across enterprise cohorts.

DAYS 1–30

Inventory and triage. List every live AI initiative. For each: the business decision or process it touches, the baseline (if any), the owner (if any), the run cost. Sort into three buckets — Evidenced, Instrumentable, Unmeasurable. Be ruthless about the third bucket; it is usually the largest.

DAYS 31–60

Instrument the middle. For the Instrumentable bucket, define one outcome metric per initiative (cycle time, error rate, rework, conversion, cost-per-transaction), capture the pre-AI baseline from historical data, and build measurement into the workflow — not into a slide. Assign a single accountable owner per initiative with kill authority.

DAYS 61–90

Decide and report. Retire or pause the Unmeasurable bucket — redeploy the spend. Take the first 30 days of instrumented readings to the board as your opening Evidenced-Value Ratio, with a target trajectory. The number will be uncomfortable. Reporting it anyway is what separates governance from theatre.



The Value Realist

Pradeepta Mishra

VP, AI Innovations, Beghou

Q1

Despite significant investment, many organisations still struggle to show a clear return on AI. From your experience, what is the real reason behind this gap - is it the technology, the measurement, or something else?

A

The biggest gap is rarely the technology. Most organisations can build models and deploy copilots today. The real issue is that AI initiatives are often disconnected from measurable business outcomes. Teams focus on accuracy, model performance, or proofs-of-concept, while business leaders care about revenue growth, cost reduction, risk mitigation, and productivity improvements.

A second challenge is that many organisations do not establish baseline metrics before deployment, making it difficult to prove causality. Successful AI programs start with a business problem, define value metrics upfront, and treat AI as a business transformation initiative rather than a technology project

Q2

If a leadership team asked you how to decide which AI initiatives to continue, scale, or stop, what advice would you give them — and what should they look at first?

A

I would ask leaders to evaluate every AI initiative across three dimensions: business value, adoption, and scalability.

1. Continue and scale initiatives that demonstrate measurable business impact, strong user adoption, and a clear path to enterprise deployment.
2. Refine initiatives that show user engagement but have not yet translated into measurable business value.
3. Stop initiatives that remain technology experiments without a credible business case or sustained adoption.

The first thing leadership should look at is not model performance—it is whether the initiative is moving a business KPI that matters. If an AI solution cannot demonstrate impact on revenue, cost, risk, customer experience, or employee productivity, it is unlikely to justify continued investment regardless of how sophisticated the technology may be.

A simple principle I often use is: **"Don't scale AI because it works; scale AI because it creates measurable business value."**

Note: CANDID features direct contributions from CXO leaders. To contribute,

contact: retina@reva.edu.in



5 Roles With Exploding Demand · Week of 1 June 2026

Focus: AI Value, Economics & Assurance

#	Role	Critical Skills	CTC Range*
1	AI Value Realisation Lead	Benefit tracking · Baseline design · Business case audit · Board reporting	₹40L–₹85L
2	AI FinOps / Cost Engineer	Token & inference economics · GPU utilisation · Cloud cost governance · Unit economics	₹30L–₹65L
3	AI Audit & Assurance Manager	EU AI Act · ISO 42001 · Model risk audit · Impact assessment	₹28L–₹60L
4	AI Product Owner (P&L-Mandate)	Outcome metrics · Workflow integration · Adoption analytics · Kill-criteria discipline	₹35L–₹70L
5	Decision Intelligence Analyst	Causal measurement · A/B & holdout design · BI storytelling · Python/SQL	₹18L–₹40L

*Indicative mid-to-senior ranges for major Indian metros, synthesised from 2026 salary reports (Taggd, upGrad, Deloitte rewards surveys). Verify against the cited reports before publication — these roles are emerging and India-specific survey data remains nascent.



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

RACE Research Radar highlights the latest publications by RACE learners, mentors, and faculty across AI, Cybersecurity, Cloud, and Analytics — applied research that moves from academic inquiry to enterprise impact.

This week, RACE researchers feature four applied AI papers with IEEE Xplore and Springer links now live. The work spans railway infrastructure intelligence, sales enablement, inclusive education, and creator-economy analytics — each converting a manual, time-intensive workflow into an AI-assisted one.

COMPUTER VISION · RAILWAY INFRASTRUCTURE · ENGINEERING AUTOMATION

Railway Track Layout Asset Reader Using Computer Vision

Ravichandran R, Prof. J. B. Simha, Dr. Shinu Abhi · ICAAIC 2025, IEEE Xplore

A computer vision system that automatically identifies critical railway assets — switches, crossings, signals, bridges, and tunnels — directly from track layout drawings, replacing time-consuming manual review. Engineering drawings are one of the last large pockets of untapped enterprise data; this work shows how vision AI converts legacy documentation into a queryable asset inventory. Directly relevant to railways, utilities, and infrastructure owners sitting on decades of paper-era schematics.

<https://ieeexplore.ieee.org/document/11330356>

SALES ENABLEMENT · RAG · SMALL LANGUAGE MODELS

Retrieval-Augmented Generation and Language Model for Advanced Product Querying

Tahiyya Khan, Dr. J. B. Simha, Paramesh G · IEEE CONECCT 2025, IEEE Xplore

An AI chatbot built on Retrieval-Augmented Generation and the Phi-3 small language model that lets sales representatives pull accurate product information from brochures, databases, and websites through a single unified interface. The choice of a compact, locally deployable model over a large API-based one is the enterprise signal here: grounded answers, lower inference cost, and data that never leaves the organisation — a pattern this week's ROI theme says more deployments should follow.

<https://ieeexplore.ieee.org/document/11306572>

EDUCATION AI · TRANSFORMERS · ACCESSIBILITY

Document Processing and Summarization with Transformers

Pavana Priya, J. B. Simha, Shinu Abhi · CML 2025, Springer

A transformer-based summarisation system that helps teachers and caregivers simplify educational content for specially abled children. Beyond its social impact, the work demonstrates a capability every enterprise L&D and communications team needs: automatically adapting complex source material to the comprehension level of a specific audience — the same architecture that simplifies a lesson can simplify a policy, a compliance manual, or a product guide.

https://link.springer.com/chapter/10.1007/978-981-95-2875-2_29

RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

CREATOR ECONOMY · LLM AGENTS · AUDIENCE ANALYTICS

Multiactivity Agent Using LLM for YouTubers

Shreekanth Suryanarayana Reddy, Krishna Kumar Tiwari, Rashmi Agarwal · ICCIDS 2025, Springer

A deployed multi-feature LLM agent for content creators that combines comment sentiment analysis, RAG-based automatic replies, a comment chatbot, and transcript summarisation in one Streamlit application. With 500 hours of video uploaded to YouTube every minute, audience engagement at scale is a data problem — and the same agent pattern applies wherever an organisation faces high-volume unstructured feedback: customer reviews, support tickets, social listening.


https://link.springer.com/chapter/10.1007/978-3-031-98360-3_17

RETINA TAKE

Four papers, one pattern: AI absorbing the manual middle of knowledge work — reading engineering drawings, answering product queries, simplifying documents, and managing audience engagement. Each converts hours of human effort into seconds of machine effort with the human retained for judgment. RACE's applied research pipeline keeps moving from theory to deployed, enterprise-relevant systems. Explore the full archive: racelit.netlify.app

BOOK SHELF

RETINA'S READ OF THE WEEK



RETINA RECOMMENDS · THIS WEEK'S THEME
All-in On AI: How Smart Companies Win Big with Artificial Intelligence

Thomas H. Davenport & Nitin Mittal · Harvard Business Review Press, 2023

Davenport — whose *Competing on Analytics* anchored Issue #3's Book Shelf — teams up with Deloitte's Nitin Mittal to study the small set of companies that stopped piloting AI and rebuilt their operating models around it: the 'all-in' enterprises. The book's enduring contribution to this week's theme is its dissection of how these organisations think about value — not

as a post-hoc ROI calculation, but as a portfolio of business outcomes with owners, baselines, and consequences. In a week when 42% of companies are abandoning AI projects they never instrumented, this is the corrective text. Every CEO and CFO preparing an AI budget review should read it before the review — and notice that not one of the companies profiled measured success in licences activated.

All-in On AI — Davenport & Mittal, HBR Press 2023:

<https://store.hbr.org/product/all-in-on-ai-how-smart-companies-win-big-with-artificial-intelligence/10520>

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

UPCOMING · JUNE 2026

Certified Agentic AI Engineer

New Batch Starting. Practical expertise in LLMs, AI Agents, Multi-Agent Systems, Agentic Workflows, Governance, and capstone-based implementation. Real-world use cases. Deployable solutions.

Apply now:

<https://race.reva.edu.in/certified-agentic-ai-engineer>

UPCOMING · SEPTEMBER 2026 BATCH
APPLY BY 31 JULY

Applications 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 now: race.reva.edu.in

COMING SOON

THE BIGGEST TECH CONCLAVE IN TOWN – RACE360 and Top Women Tech Awards

<https://race.reva.edu.in/race360>

<https://race.reva.edu.in/top-10-women-tech-leaders-india-awards>

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 LEADERSHIP TEAM

“If we cancelled every AI initiative we cannot measure tomorrow morning - how much budget would we recover, and what, honestly, would we lose?”

RETINA PUBLISHES EVERY SUNDAY · FORWARD TO A CXO YOU RESPECT