

RACE WEEKLY INTELLIGENCE BRIEF


RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER:

THE END OF PER-SEAT SOFTWARE

For twenty years, software was sold by the seat and IT services by the hour. Both were really charging for one thing: human effort. AI just broke that. When an agent resolves a support ticket for 99 cents, or does the work of ten people, nobody is buying ten seats. Seat-only pricing fell from 21 to 15 percent of SaaS firms in a single year. In India, TCS now prices about 80 percent of its finance and HR contracts on outcomes, not hours. This changes your software bill, your services contracts, and if you sell either, your business. Here is what is happening, and what to do about it.



ALSO INSIDE

The world in 5 · Signal · Hyperfocal ·
The new receipt · Perspectives · The
repricing playbook · Metric of the
week · Fun fact · Candid · Book shelf
· Research radar · RACE pulse

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

FROM THE EDITOR'S DESK

You never really paid for software. You paid for the person using it. That quiet assumption just stopped holding.

Per-seat pricing made perfect sense for twenty years. A seat was a person, a person did work, so charging per seat was a fair proxy for the value delivered. The same logic ran under IT services, only there the proxy was the hour. Both models priced human effort, dressed up as a licence or a timesheet. The whole industry, software and services alike, was in the business of renting out people, or access to the tools people used.

Agents break the proxy. When the software does the work itself, there is no seat to count and no hour to bill. A vendor cannot charge ten licences for a job one agent now does, and a client will not pay for a hundred consultants when fifty and an agent deliver the same result. So pricing is moving to the only thing left that still carries value: the outcome. Intercom charges 99 cents per resolved conversation. Salesforce charges per action and per conversation. And in India, TCS says about 80 percent of its finance and HR contracts are now tied to outcomes, roughly double the share two years ago.

Two honest notes keep this grounded. The seat is not vanishing overnight. Seat-only pricing is losing its default status, not disappearing, and the common contract now has three parts: a platform fee, some human licences, and a pool of agent activity. And the shift is uneven. Outcome pricing is winning first where results are frequent and easy to verify, like customer support, and slowly elsewhere. But the direction is set, and it reaches every leader in two ways at once. It changes what you pay for the software and services you buy. And if you sell either, it changes what you are allowed to charge. This issue maps both sides.

THE WORLD IN 5

GLOBAL HEADLINES A CXO SHOULD BE ABLE TO REFERENCE

- From The Editor's Desk
- **The World In 5**
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

Five signals from the past fortnight, worth a note at your next leadership table.

1. THE SEAT IS LOSING ITS DEFAULT STATUS

Seat-based pricing fell from 21 percent of SaaS firms to 15 percent in a single year, while hybrid models rose from 27 to 41 percent. The reason is simple. An agent does the work and never logs in, so there is no seat to charge for. Vendors are moving to usage, outcomes, and blends of the two. Per seat is not dead, but it is no longer the obvious default.

2. OUTCOME PRICING IS ALREADY LIVE

This is not a forecast. Intercom's Fin charges 99 cents only when it resolves a conversation, and that one move took the product from about 1 million to over 100 million dollars in annual revenue. Salesforce now lists four meters on one pricing page: seats, actions at 10 cents, conversations at 2 dollars, and resolutions. Zendesk and HubSpot charge per resolution too. You pay when it works.

3. THE REVENUE MATH FLIPS

When software takes on the labour rather than just providing access, venture investors at a16z estimate it can earn 2 to 10 times more revenue per customer. The flip side is pressure. Bain warns that vendors who do not move their pricing within about 18 months face permanent revenue erosion. The model rewards those who price to the work and punishes those who cling to the seat.

4. BUYERS ARE ALREADY SAVING

One Fortune 500 financial firm cut its software seat count by 40 percent after deploying AI agents in customer service, saving about 12 million dollars a year. Expect more of this at every renewal. The leverage in the room has shifted toward the buyer who can say, plainly, that they no longer need the seats they used to.

5. INDIA: THE EFFORT MODEL UNDER REAL PRESSURE

India's 315 billion dollar IT services industry was built on billing for people and hours. That link is now stretching. TCS says about 80 percent of its finance and HR contracts are outcome-based, and it cut more than 23,000 roles in FY26 in a pivot to AI-first delivery. Its chairman expects the firm to soon run as many AI agents as it has employees. For India, this is the single biggest commercial shift in a generation.

Source:

<https://www.growthunhinged.com/p/the-state-of-b2b-mon-etization-in-2026>

<https://www.intercom.com/pricing>

<https://www.salesforce.com/in/agentforce/pricing/>

<https://a16z.com/vertical-saas-now-with-ai-inside/>

<https://www.reuters.com/world/india/ai-reshapes-indias-it-services-sector-contracts-clients-demand-more-less-2026-08-20/>

 SIGNAL

THE REPRICING, BY THE NUMBERS

- From The Editor's Desk
- The World In 5
- **Signal**
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



15x

of SaaS firms still price mainly by the seat, down from 21% a year earlier; hybrid rose from 27% to 41% (Growth Unhinged)

\$0.99

Intercom Fin's price per resolved support conversation, the move that took it from about \$1M to \$100M+ in revenue

80%

of TCS finance and HR contracts now outcome-based, about double the share since late 2023 (Reuters 2026)

2 to 10x

more revenue per customer when software takes on the labour, not just access (a16z estimate)

\$12M

yearly saving at one Fortune 500 firm that cut software seats 40% after deploying AI agents

18 mo

the window Bain gives seat-only vendors to reprice before facing permanent revenue erosion

Source:

<https://www.growthunhinged.com/p/2025-state-of-b2b-monetization>

<https://www.intercom.com/pricing>

<https://www.reuters.com/world/india/ai-reshapes-indias-it-services-sector-contracts-clients-demand-more-less-2026-08-20/>

<https://a16z.com/vertical-saas-now-with-ai-inside/>



HYPERFOCAL

WHAT IS REALLY GOING ON

- From The Editor's Desk
- The World In 5
- Signal
- **Hyperfocal**
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

The seat was always a proxy. The proxy just broke.

Think about what a software seat actually measured. It was never the value of the software by itself. It was a stand-in for a person who logged in and did something useful with it. One seat, one worker, one stream of value, and a price that tracked all three together. Services pricing did the same job with a different ruler. The billable hour was a proxy for a person applying effort. For two decades, both models worked because the person and the value moved together. Where there was value being created, there was a human creating it, and you could price the human.

Agents separate the two. Now the value can be created without a person in the seat and without an hour on the clock. The proxy that held the whole pricing model together simply comes apart. A vendor that keeps charging per seat is charging for a thing that no longer maps to value, which means either the customer overpays for empty seats or the vendor watches its revenue fall as seats disappear.

Neither is stable. So pricing is migrating to the one thing that still carries the value directly: the result. Pay when the ticket is resolved. Pay when the action completes. Pay for the outcome delivered.

For a leader, this is not a software procurement footnote. It is a shift in what money buys. On the buying side, your costs stop tracking your headcount and start tracking your usage and your results, which is better in good times and sharper in bad ones. On the selling side, if your company bills for seats or hours, a rival who bills for outcomes is quietly teaching your customers what your work should cost, and usually that it should cost less. The honest nuance holds. Seats will not vanish, and outcome pricing fits some work far better than others. But the question every leader should now ask, as buyer and as seller, is the same. What am I actually paying for, and what am I actually charging for, now that the seat no longer stands for the work.

RETINA TAKE

The seat and the hour were always proxies for human effort. AI severed effort from the human, so the proxies stopped tracking value and began tracking nothing. Whoever prices the outcome wins the next decade. Whoever keeps counting seats is counting something that no longer means what it used to.

THE NEW RECEIPT

WHAT THE BILL USED TO SAY, AND WHAT IT SAYS NOW

A new segment for this issue. Three kinds of work, shown as the old line item and the new one, with what the change means for a buyer. The examples are real and current.

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- **The New Receipt**
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

Customer support

FROM A SEAT PER AGENT TO 99 CENTS PER ANSWER

Old receipt: a monthly licence for every support agent, whether the queue was busy or empty. New receipt: Intercom's Fin bills 99 cents when it resolves a conversation, and nothing when it fails and hands off to a human. Zendesk and HubSpot price the same way, per resolved ticket. The cost now rises and falls with the work, not the roster.

What it means: Your support software bill should now move with volume and success, not headcount. If it still tracks seats, you are paying for chairs no one sits in.

Sales and operations

FROM A SEAT PER REP TO A CHARGE PER ACTION

Old receipt: a per-user licence for the platform, plus more users as you grew. New receipt: Salesforce Agentforce meters the agent's work directly, about 10 cents per standard action or 2 dollars per conversation, sitting on top of, not instead of, the seats. One pricing page, four meters. The trap is reading only the one that fell.

What it means: When a vendor cuts the seat price and adds consumption meters, model the whole bill under heavy agent use before you sign. The headline can drop while the total climbs.

Software modernization

FROM A TEAM FOR A YEAR TO AN OUTCOME FOR 60 PERCENT LESS

Old receipt: a large offshore team, billed by the hour over many months, to migrate a legacy system. New receipt: Infosys reported moving 3 million lines of a client's old code to a modern stack using AI, at about 60 percent lower cost and on a 60 percent shorter timeline. Once a vendor advertises that, it has taught every future buyer what the job should cost.

What it means: For large effort-based projects, ask vendors to price the outcome, not the team. The first one who will is setting the new market rate for everyone.

PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- **Perspectives**
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

WE BUILT AN INDUSTRY ON BILLABLE HOURS. THE CLOCK IS BEING UNPLUGGED.



Dr. Shinu Abhi
Professor and Director
Corporate Training, RACE,
REVA University

<https://www.linkedin.com/in/shinu-abhi-phd-711a339/>

I live in the middle of this story. We are in Bengaluru, we train thousands of working professionals, and a large share of them come from an industry whose entire commercial model is the billable hour. For a generation, India's IT services sector turned talent into revenue in a beautifully simple way: hire people, train them, bill their time. More people meant more revenue. That link held for so long that it stopped feeling like a choice and started feeling like a law of nature. It was never a law. It was a pricing model, and pricing models end.

The numbers are now impossible to look past. TCS says about 80 percent of its finance and HR contracts are outcome-based, it reduced headcount by more than 23,000 in a single year, and its chairman has said the firm will soon run as many AI agents as it employs people. Read that last sentence slowly. The most valuable worker in the next contract may not draw a salary. When a client can be shown an outcome delivered for 60 percent less, the old bill, priced by the hour, simply stops being credible. This is not a downturn you wait out. It is a repricing you adapt to.

Here is the part I find hopeful, and it is why I do the work I do. The demand for the outcome has not fallen. If anything it has risen. What changes is where the human value sits. It moves from doing the repetitive work to designing, directing, and governing the agents that now do it, and from billing

hours to owning results. That is a higher-value place to stand, but only for people who make the shift deliberately. For India this is the defining reskilling task of the decade. We teach it as frugal, human-first AI, and the message to every professional and every firm is the same. Stop selling the hour. Start owning the outcome.

RACE Perspective

The billable hour and the software seat were the same idea wearing two uniforms, both pricing human effort. AI retired the idea, not the value. The winners will move their people, and their pricing, from effort to outcome on purpose, before a competitor or a client does it for them.

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- **The Repricing Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE REPRICING PLAYBOOK

WHAT TO DO AS A BUYER,
AND AS A SELLER

Five moves, ordered from fastest to most strategic, that any leadership team can act on.

MOVE 1

AUDIT WHAT YOU ARE ACTUALLY PAYING FOR

Pull your top software and services contracts and ask one question of each line. Are we paying for seats or hours that agents have started to replace. The Fortune 500 firm that saved 12 million dollars did not get a discount. It stopped paying for seats it no longer needed. The same slack is almost certainly sitting in your contracts, waiting for the next renewal.

The move: Review your largest contracts for seats and hours displaced by AI, and take the savings to the table at renewal. Nobody will offer it to you.

MOVE 2

WATCH THE TOTAL BILL, NOT THE HEADLINE PRICE

The common vendor play is to cut the seat price and quietly add consumption and outcome meters on top. One pricing page, several meters. If you read only the meter that fell, your total cost can rise while you congratulate yourself on the saving. Model your spend under heavy agent use, across every meter, before you sign anything.

The move: Before you sign a repriced contract, model the full bill under real agent volume. A lower seat price with four new meters is often a bigger invoice.

MOVE 3

RE-SKILL PROCUREMENT FOR USAGE AND OUTCOMES

Buying by the seat was predictable. You counted people and multiplied. Buying by usage and outcome is a different discipline, closer to buying cloud than buying licences, with forecasting, unit economics, and the risk of a surprise bill. Finance and procurement need that muscle now, not after the first shock invoice arrives.

The move: Build the capability to forecast and negotiate consumption and outcome contracts this quarter. It is a new skill, and the vendors already have it.

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- **The Repricing Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE REPRICING PLAYBOOK

WHAT TO DO AS A BUYER,
AND AS A SELLER

MOVE 4

IF YOU SELL, PRICE TO VALUE, NOT TO EFFORT

If your company bills for seats or hours, a competitor billing for outcomes is teaching your customers what your work should cost, and usually that it should cost less. Waiting is the one move that guarantees you lose the comparison. Find a single offering where you can measure the result, and pilot pricing it on that outcome, with a floor to protect you while you learn.

The move: Pick one offering you can tie to a measurable result and pilot outcome pricing on it now. The first mover in your category sets the reference price.

MOVE 5

DELINK YOUR REVENUE FROM HEADCOUNT ON PURPOSE

For services firms, and India's especially, the link between headcount and revenue is breaking whether you plan for it or not. The deliberate path is to move your best people from doing the repetitive work to designing and governing the agents that do it, and to price the result. The accidental path is to keep billing hours until a client shows you the 60-percent-cheaper quote.

The move: Shift people from doing the work to governing the agents that do it, and reprice around outcomes. Choose the move before the market makes it for you.

Source:

<https://www.growthunhinged.com/p/the-state-of-b2b-monetization-in-2026>

<https://a16z.com/vertical-saas-now-with-ai-inside/>

<https://www.reuters.com/world/india/ai-reshapes-indias-it-services-sector-contracts-clients-demand-more-less-2026-08-20/>

<https://www.infosys.com/investors/reports-filings/annual-report/annual-reports/ar-2025-26.html>

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- **Metric Of The Week**
- **Fun Fact**
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

METRIC OF THE WEEK

ONE NUMBER, PROPERLY UNDERSTOOD

4 meters on a single software pricing page. The era of one price per seat is over.

What it is: Salesforce's Agentforce pricing now carries four different meters side by side: a per-user seat at 125 dollars a month, actions at about 10 cents each, conversations at 2 dollars, and per-resolution pricing for its help agent. One product, four ways to be billed. A decade ago there was one number on that page, the price of a seat.

Why it matters now: The four meters are the whole shift in miniature. The seat did not disappear, it became one line among several, and most of the new value is being metered somewhere else. A buyer who still negotiates only the seat price is negotiating the one number that matters least.

How leaders use it: When you open a vendor's pricing, count the meters. Then

ask which one carries most of your real usage, and model your bill there. The answer is rarely the seat anymore, and the vendor is counting on you not to check.

Source:
<https://www.salesforce.com/in/agentforce/pricing/>

FUN FACT

SOMETHING TO REPEAT AT DINNER

A software vendor is now so sure its product works that it will pay you if it does not. Intercom's Fin resolves customer conversations for 99 cents each, and the company has offered to refund up to 1 million dollars if it misses agreed resolution targets. Sit with how strange that is. For two decades, software was sold on the promise that it might help, and the risk of it not working sat entirely with you, the buyer. Now the vendor is putting a million dollars behind the result. When the seller starts carrying the risk of the outcome, you are no longer buying a tool. You are buying a result, and the price of just having access is quietly heading toward zero.

Source: <https://www.intercom.com/fin/pricing> | <https://www.intercom.com/help/en/articles/11643174-what-is-the-fin-guarantee>

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- **Candid**
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



CANDID



Sandeep Vijayaraghavan

EVP – Cyber Security & Cloud Services, Terralogic

<https://www.linkedin.com/in/sandeep-vijayaraghavan/>

Q1. Seat-only pricing fell from 21 to 15 percent of SaaS firms in a year. Is per-seat really dying, or just losing its default status, and how should a buyer read the difference?

Ans: No, but it is no longer the default. CXOs should price technology on usage, business value and scalability, not headcount alone.

Q2. Vendors are cutting seat prices while adding consumption and outcome meters. How should a CXO make sure the total bill does not quietly rise even as the headline falls?

Ans: Control total cost of ownership. Scrutinise consumption charges, add-ons and escalation clauses; mandate transparent pricing, define usage limits and review consumption continuously.

Q3. TCS now prices about 80 percent of some contracts on outcomes rather than hours. For India's IT services industry, is outcome pricing a growth opportunity or a margin trap?

Ans: It is a major growth opportunity when outcomes are measurable, accountability is explicit and risk is disciplined. Vague targets and unrealistic commitments will destroy margins.

Q4. If you sell software or services today, what is the first thing that has to change to move from charging for effort to charging for results?

Ans: Shift from measuring effort to measuring business impact. Define precise KPIs, assign ownership and align delivery commitments with outcomes the customer values.

Note: CANDID features direct contributions from CXO leaders.
To contribute, contact: retina@reva.edu.in



BOOK SHELF

RETINA'S READ OF THE WEEK

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- **Book Shelf**
- Research Radar
- RACE Pulse
- RACE Programs

RETINA RECOMMENDS · THIS WEEK'S THEME


Monetizing Innovation: How Smart Companies Design the Product Around the Price *Madhavan Ramanujam and Georg Tacke · Wiley, 2016*

Ramanujam and Tacke, from the pricing firm Simon-Kucher, argue that most companies build a product first and think about price last, which is exactly backwards. Their case is that price is not what you decide at the end. It is the clearest signal of the value you are creating, and you should design around it from the start. Read against this week's cover, the book is almost eerily on point. The move from seats to outcomes is precisely a move from pricing access to pricing value, which is the whole argument of the book made real across an entire industry at once.

What it gives a leadership team is a vocabulary for the conversation every

board is now having. Willingness to pay, value metrics, good-better-best, and how to tie a price to the outcome a customer actually cares about. If you sell anything, it is a practical guide to the repricing this issue describes. If you buy, it is a map of the tactics about to be used on you.

It predates the agent era, which is its strength. The principles are not about AI, so they survive it. The seat and the hour were simply value metrics that worked for a while and now do not. This book is about choosing the next one on purpose. Read it before your next pricing or procurement review.



- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- **Research Radar**
- RACE Pulse
- RACE Programs



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

Cybersecurity · vulnerability management · semantic learning

Context-Aware Enterprise Vulnerability Prioritization and Remediation Recommendations Using Semantic Learning

Sumit Prakash Dubey; Soumya Maity

Every security team faces the same problem: far more flagged vulnerabilities than it can ever fix, and a severity score that says nothing about what matters to your business. This work moves the decision from generic severity to real organisational risk. It reads vulnerability and remediation text with a Sentence-BERT model for meaning-based matching, adds a risk model that weighs asset importance, exposure, and likelihood of exploitation, then re-ranks the fixes with a learning-to-rank approach, so analysts get a clear, ordered list of what to remediate first. It beat lexical, semantic-only, and heuristic methods on ranking quality while keeping queries fast, and it shows its working through source attribution and decomposed scores.

Enterprise relevance: This is how a security team spends a fixed budget on the right fixes. Ranking by business risk rather than raw CVSS means scarce analyst hours land on the vulnerabilities that actually threaten critical assets, and the built-in explainability and audit trail is exactly what a CISO needs to defend the decision to a board or a regulator.

2026 5th International Conference on Sentiment Analysis and Deep Learning (ICSADL) · IEEE

<https://ieeexplore.ieee.org/document/11451870>

Cybersecurity · threat intelligence · NLP

Threat Watch Pro: NextGen Unified Threat Intelligence Platform

Sudharsan Balaji; Govardhan Sridhar; Shinu Abhi; Rashmi Agarwal

Threat intelligence arrives from too many sources, in too many formats, to act on quickly, which leaves teams drowning in data and prone to false positives and missed signals. This open-source platform pulls it into one place and maps it to the MITRE ATT&CK and D3FEND frameworks, turning raw feeds into structured, actionable intelligence. A natural-language layer lets both technical and non-technical staff query and interpret threats without specialist skills, and the authors report automated threat mapping lifting response efficiency by over 45 percent through streamlined reporting and cross-team collaboration.

Enterprise relevance: Fragmented threat intelligence is slow and expensive, and slowness is where breaches live. Consolidating feeds, mapping them to a shared framework, and letting any stakeholder ask questions in plain language turns intelligence from a specialist bottleneck into a shared, faster capability, which is precisely what shortens the window between a threat appearing and a response.

Business Intelligence and Data Analytics (BIDA 2025) · Smart Innovation, Systems and Technologies, vol 472, pp 275–287 · Springer, Cham, 2026

https://link.springer.com/chapter/10.1007/978-3-032-15996-0_21

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

- From The Editor's Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

Built for the repricing: Certified Agentic AI Engineer APPLY BY 31 OCTOBER 2026

If this week's theme has a practical answer, it is people who can design and govern the agents that are starting to do the work. RACE's Certified Agentic AI Engineer programme is built for exactly that: LLMs, AI agents, multi-agent systems, agentic workflows, governance, and a capstone that ships a deployable agent. For professionals whose roles are moving from doing the work to directing it, this is the most direct reskilling path we run.

To know more, call:
+91 89040 58866

enquiry@race.reva.edu.in ·
Apply now: race.reva.edu.in

Master's and PG programmes APPLY BY 31 OCTOBER 2026

Applications are open across RACE's Master's and PG programmes: Business Analytics (ranked #1 in India), Artificial Intelligence (#3), Cybersecurity (#5), and Cloud Architecture and Security. Weekend format, AICTE and UGC approved, built for working professionals who want to move up the value chain rather than be repriced down it. Details at race.reva.edu.in.

To know more, call:
+91 89040 58866

enquiry@race.reva.edu.in ·
Apply now:
<https://race.reva.edu.in/certified-agenticai-engineer>

RACE Sandbox and CXO Circle SATURDAY 31 OCTOBER 2026

RACE Sandbox, the on-campus career-simulation open day, returns through the term: sit in a live class, walk the AI, Cybersecurity and Cloud labs, and take a 1:1 on where your career goes next. For senior leaders, the CXO Circle breakfast continues as a closed-door peer conversation on exactly the shifts RETiNA covers.

9:30 AM – 2:00 PM
RACE, REVA University,
Bengaluru

Reserve your slot:
+91 89040 58866 ·

- From The Editor’s Desk
- The World In 5
- Signal
- Hyperfocal
- The New Receipt
- Perspectives
- The Repricing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- **RACE Programs**

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

PG Diploma / M.Sc. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

12–24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

PG Diploma / M.Sc. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic.

12–24 Months

#3 IN INDIA · TECHGIG

ARTIFICIAL INTELLIGENCE

M.Tech. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

M.Tech. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic

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

AGENTIC AI | 4–6 weeks

AGENTIC AI SYSTEMS DESIGN

Multi-agent systems using LangChain, LangGraph, CrewAI. Capstone:
deployed autonomous agent.

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.

CYBERSECURITY | 8–12 WEEKS

CYBERSECURITY & SOC OPERATIONS

Threat detection, incident response, SIEM/SOAR. Hands-on in RACE's live
Security Operations Centre.

DATA & AI | 12 WEEKS

DATA SCIENCE & BUSINESS ANALYTICS

Statistics to ML to BI. Python, SQL, Power BI, Tableau on real business
datasets.

LEADERSHIP | 8–12 WEEKS

TECHNO-FUNCTIONAL LEADERSHIP

Executive presence, decision-making, leading tech-driven teams. Built for
mid-to-senior managers.

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



REVA Academy for
Corporate Excellence

RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK, ASK YOUR LEADERSHIP TEAM

“On everything we buy and everything we sell, are we still pricing by the seat and the hour? And if a competitor priced the outcome instead, who would our customers call?”

REVA Academy for Corporate Excellence,
REVA University
Rukmini Educational Charitable Trust
Kattigenahalli, Yelahanka, Bangalore, Karnataka, India, 560 064
Phone: +91 89040 58866

RETINA PUBLISHES EVERY SUNDAY · FORWARD TO A CXO YOU RESPECT



REVA
UNIVERSITY
Bengaluru, India



Visit our website
race.reva.edu.in