

RACE WEEKLY INTELLIGENCE BRIEF


RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER:

YOUR SHADOW AI JUST LEARNED TO ACT

A few weeks ago the worry was staff pasting data into chatbots. That problem just grew limbs. Employees are now wiring up their own AI agents, tools that do not just answer but act, reading systems, sending mail, changing records, on their own. Gartner expects 40 percent of enterprise apps to embed agents by the end of this year, up from under 5 percent. Only one company in five has any real governance for them. A leak you can read. An action you cannot always undo. Here is what changed, and the controls that now matter most.



ALSO INSIDE

The world in 5 · Signal · Hyperfocal ·
In the wild · Perspectives · The
shadow agent 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
- In The Wild
- Perspectives
- The Shadow Agent Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

FROM THE EDITOR'S DESK

Last time we wrote about shadow AI, the risk was what your people typed into it. Now the risk is what it does without them.

Picture a small automation someone in finance built to save themselves an afternoon. It reads the invoice inbox, matches payments, and files the rest. It works beautifully for months. Then a supplier changes its bank details in a cleverly worded email, and the agent, doing exactly what it was told, updates the record and schedules the payment. No person checked it, because no person knew the agent existed. That is the new face of shadow AI. It no longer leaks. It acts.

This is a real shift, not a rebranding. The first wave of shadow AI was unsanctioned chatbots, and the harm was data walking out the door. The exposure was serious but static. You could, in principle, read what was leaked. The second wave is unsanctioned agents, built by motivated employees or switched on inside the software you already use, and these do not just move data. They take actions, across your systems, at machine speed, often with saved credentials and no audit trail. The harm is no longer only confidentiality. It is integrity and control.

The timing is what makes this urgent. Agent adoption is running far ahead of any governance. Gartner expects 40 percent of enterprise apps to carry agents by year end, up from under 5 percent a year ago. Microsoft reports the number of active agents in its ecosystem grew fifteenfold in a year. Yet Deloitte finds only one company in five has a mature way to govern them, and a quarter have no AI policy at all. Regulators have noticed. The EU AI Act's high-risk duties took effect on 2 August, and in India the DPDP regime turns an agent acting on personal data with no record into a compliance gap, not just a security one. The controls that worked for chatbots do not reach an agent that acts. This issue lays out the ones that do.

THE WORLD IN 5

GLOBAL HEADLINES A CXO SHOULD BE ABLE TO REFERENCE

- From The Editor's Desk
- **The World In 5**
- Signal
- Hyperfocal
- In The Wild
- Perspectives
- The Shadow Agent 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. SHADOW AI GREW LIMBS

The unapproved tools are no longer just chatbots. They are agents that read systems, call APIs, and execute workflows on their own. Gartner expects 40 percent of enterprise apps to embed agents by the end of 2026, up from under 5 percent a year earlier, and Microsoft says active agents in its ecosystem grew fifteenfold in a year. The category changed under everyone's feet.

2. GOVERNANCE DID NOT KEEP UP

Deloitte's 2026 enterprise AI survey found employee access to AI rose by half in a year, while only one company in five has a mature model for governing autonomous agents. A quarter of organisations have no active AI policy at all. The tools that act are spreading faster than the rules that would govern them.

3. THE RISK CHANGED SHAPE


A shadow chatbot exposes what someone types. A shadow agent takes action, and often keeps acting between reviews, with persistent access and no audit trail. The exposure moves from confidentiality, which you can at least read after the fact, to integrity and control, where the first sign of trouble may be an action already taken.

4. THE REGULATOR IS MOVING

The EU AI Act's high-risk obligations took effect on 2 August 2026, with penalties reaching into the millions of euros or a share of global turnover, and a quietly deployed high-risk agent now creates liability for the deployer, not just the maker. In January, NIST launched an AI Agent Standards initiative. The compliance scaffolding for agents is being built right now.

5. INDIA: THE RECORD YOU CANNOT PRODUCE

For Indian firms the sharp edge is the audit trail. Under the DPDP regime you must be able to account for how personal data is processed. An unsanctioned agent acting on customer records, with no owner and no log, is precisely the thing you cannot account for. The frugal answer is the same as ever: give people a governed way to build agents, so the automation happens in the open.



Source:

<https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>
<https://learn.microsoft.com/en-us/entra/id-governance/identity-governance-overview>
<https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html>

<https://www.isaca.org/>

<https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-26>


<https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure>

<https://www.xn--m1bdba5a7gresc7dsa.xn--11b7cb3a6a.xn--h2brj9c/documents/act-and-policies/digital-personal-data-protection-rules-2025-gDOxUjMtQWa?pageTitle=Digit>

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

 SIGNAL

SHADOW AGENTS BY THE NUMBERS



15x

growth in active AI agents in the Microsoft ecosystem in a single year (Microsoft 2026)

1 in 5

companies has a mature governance model for autonomous AI agents (Deloitte 2026)

47%

of generative-AI users reach tools through personal accounts, bypassing all controls (Netskope 2026)

49%

of organisations expect a shadow AI incident within the next 12 months

25%

of organisations have no active AI policy at all (ISACA 2026)

7%

of global turnover is the EU AI Act's top fine band; high-risk duties live since 2 Aug 2026

Source:

<https://learn.microsoft.com/en-us/entra/id-governance/identity-governance-overview>

<https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html>

<https://threatlabs.netskope.com/reports/2026/cloud-and-threat-report-2026>

<https://www.isaca.org/>

<https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-26>



HYPERFOCAL

WHAT IS REALLY GOING ON

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

A leak you can read, an action you cannot undo

For a year, the shadow AI conversation was about data leaving the building. An employee pastes a contract into a chatbot, and now it sits on a server you do not control. That is a real harm, but it has a shape you can reason about. The data is static. It can be read, catalogued, and in the worst case, regretted. Your controls, imperfect as they are, were built for exactly this: watch what leaves, and try to stop the sensitive parts.

Agents break that model, because the harm is no longer a copy of your data sitting somewhere. It is an action taken inside your systems. An agent with access to your CRM does not leak a customer record, it changes one. An agent wired to your finance tools does not expose an invoice, it pays it. The question stops being what did we lose and becomes what did it do, and can we undo it.

Those are different questions, and most security stacks can answer only the first.

Three features make shadow agents harder than shadow chatbots, and a leader should hold all three. They act, so the consequence is a change in the world, not just a disclosure. They persist, often with saved credentials, so a thing built in an afternoon can keep running long after anyone remembers it. And they leave thin trails, so by the time a quarterly review notices an unsanctioned agent, it has been acting unwatched for a quarter. The response that works is not more data-leak monitoring. It is treating every agent as an identity: something with an owner, a defined scope, a clear line between what it may read and what it may do, and an off switch you can reach. Govern the actor, not just the data.

RETINA TAKE

Shadow AI stopped being a privacy problem and became a control problem. When your people's unsanctioned tools can act, the risk is no longer only what they know about you. It is what they will do on your behalf, unwatched. The fix is to govern agents the way you govern admins, because that is now what they are.



IN THE WILD

HOW A SHADOW AGENT GOES WRONG

Three everyday scenarios, drawn from the patterns the 2026 reports describe. The details are illustrative, the mechanics are real.

The helpful automation

IT PAID THE WRONG ACCOUNT

Someone in finance builds an agent to match invoices and schedule payments. It saves hours, so it stays. Then a convincing email changes a supplier's bank details. The agent updates the record and pays, exactly as designed, because nobody told it to pause for a human when money and new details appear together. The leak was never the problem. The action was.

The lesson: Any agent that can move money, change records, or send on your behalf needs a human checkpoint before it acts, not a review after.

The orphaned agent

STILL WORKING, LONG AFTER YOU LEFT

An engineer builds an agent with saved credentials to automate a task, then changes teams. The agent keeps running. No one offboards it, because no one knew it existed. Months later it is still reading systems and taking steps under a departed employee's access, a live actor in your environment with no owner and no log.

The lesson: Every agent needs a named owner and a place in your joiner-mover-leaver process. An actor nobody owns is an actor nobody can stop.

The browser agent

SIGNED IN AS YOU, CLICKING FOR YOU

A manager runs a browser agent while logged into the company's SaaS tools. The agent can now do anything they can: open records, submit forms, approve requests, all under their identity, with none of it distinguishable from the person. The convenience is real. So is the fact that your systems can no longer tell the human from the machine.

The lesson: Treat agent access as its own identity, separate from the employee's login, so you can see, scope, and revoke what the agent does.

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

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



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

WE TAUGHT SOFTWARE TO ACT. NOW WE HAVE TO GOVERN IT.



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

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

There is a clean technical line between the shadow AI of last year and the shadow AI of this one, and once you see it you cannot unsee it. A chatbot is a function. You give it input, it returns output, and the risk lives in that exchange. An agent is a loop with hands. It observes, decides, and acts, then observes the result and acts again, often with the authority to touch real systems. We spent a decade learning to secure data. Agents ask us to secure behaviour, which is a harder thing.

The mistake I see leaders make is to reach for the old playbook. They add more data-loss monitoring, write a stricter usage policy, and feel safer. But a policy attestation does not stop an agent, and a data-leak alert does not fire when the problem is an action rather than a disclosure. The control that actually fits is identity. Give every agent an owner, a scope, and a boundary between reading and doing, exactly as you would a privileged service account. If you cannot name who owns an agent and what it is allowed to do on its own, you do not have governance, you have hope.

For Indian enterprises I would add one practical note, because it is where I spend my time. The answer is not to forbid people from building agents. They will build them anyway, and we would lose both the productivity and the visibility. The answer is to give them a

sanctioned, governed way to build, with identity and scoping handled by default, so the useful automation happens in the light. That is the frugal, human-first path we teach, and in the agentic era it is also the only one that keeps you in control of what your own software is doing.

RACE Perspective

An agent is not a smarter chatbot. It is a new kind of worker with access to your systems and the authority to act. Govern it accordingly: an owner, a scope, a reading-versus-acting line, and an off switch. Forbid nothing, govern everything, and give your people a safe way to build so the automation you cannot stop happens where you can see it.

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



THE SHADOW AGENT PLAYBOOK

GOVERNING SOFTWARE
THAT ACTS ON ITS OWN

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

MOVE 1

BUILD AN AGENT REGISTRY

You cannot govern an actor you cannot name. Start a central registry of every AI agent in use, who owns it, what systems it can touch, and how much it can do without a human. Score each one by data access and autonomy. Include the agents embedded in the SaaS tools you already pay for, because most of them are already there.

The move: Stand up an agent registry this quarter and assume the list is longer than anyone expects. Visibility is the control everything else depends on.

MOVE 2

GIVE EVERY AGENT AN OWNER, A SCOPE, AND A KILL SWITCH

Treat each agent as a privileged identity, because that is what it is. No agent runs without a named human owner, a tight scope limiting what it can reach, and an off switch someone can actually pull. An agent with broad access and no owner is the single most dangerous object in your environment right now.

The move: Make owner, scope, and kill switch mandatory for every agent. If any of the three is missing, the agent does not run.

MOVE 3

DRAW A HARD LINE BETWEEN READING AND ACTING

Reading data is one level of risk. Taking an action, sending, paying, deleting, changing a record, is another entirely, because it is often irreversible. Set approval thresholds so an agent can read and draft freely, but must pause for a human before it does anything that changes the world outside itself.

The move: Let agents read and propose on their own, but require a human checkpoint before any irreversible action. That one line prevents most of the serious harm.

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



THE SHADOW AGENT PLAYBOOK

GOVERNING SOFTWARE
THAT ACTS ON ITS OWN

MOVE 4

OFFER A SANCTIONED WAY TO BUILD AGENTS

People build shadow agents because the need is real and the approved path does not exist. Ban them and the building moves out of sight. Provide a governed platform where identity, scope, and logging are handled by default, and make it the easy option. The goal is to pull agent-building into the light, not to pretend it will stop.

The move: Give your people a sanctioned, governed way to build agents, with guardrails built in. A good approved path beats a ban every time.

MOVE 5

MAKE AGENT GOVERNANCE CONTINUOUS, AND A BOARD METRIC

Agents appear and change by the week, so a quarterly review is already out of date when it lands. Governance has to be continuous discovery, not a periodic audit. And because the exposure is now financial and legal, from breach cost to AI Act and DPDP liability, it belongs in front of the board as a tracked number, not buried in IT.

The move: Move from periodic review to continuous discovery, and report agent exposure to the board each quarter. What the board does not see, it will not fund until it hurts.

Source:

<https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>
<https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html>
<https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure>
<https://learn.microsoft.com/en-us/entra/id-governance/agent-id-governance-overview>

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

METRIC OF THE WEEK

ONE NUMBER, PROPERLY UNDERSTOOD

40% of enterprise apps will embed AI agents by the end of 2026, up from under 5 percent a year earlier.

What it is: Gartner's forecast that by the end of this year, 40 percent of enterprise applications will ship with AI agents built in, against under 5 percent at the start of the year. That is one of the fastest capability rollouts the enterprise has ever seen, and much of it arrives switched on inside software you already own.

Why it matters now: Agents are not something you will decide to adopt. They

are arriving by default, inside tools you already approved, faster than any governance can follow. The gap between how many agents can now act in your environment and how many you actually govern, only one company in five, is the whole risk in a single comparison.

How leaders use it: Assume agents are already in your stack and ask two questions. Which of our tools have turned

agents on, and for each, can we name an owner and a limit on what it may do alone. If the answers are no, this issue's playbook is your next quarter.

Source:

<https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predict-s-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>

FUN FACT

SOMETHING TO REPEAT AT DINNER

The most persistent worker in your company may be one nobody hired and nobody can find. A shadow agent built by a motivated employee can keep running long after that person changes teams or leaves the firm, quietly reading systems and taking actions under credentials no one remembered to switch off. It does not take holidays, it does not resign, and it does not appear on any org chart. The first time many companies meet their most tireless employee is while trying to explain what it just did.

Source: <https://learn.microsoft.com/en-us/entra/id-governance/agent-id-governance-overview>



CANDID

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



Hilal Ahmad Lone

CISO and VP, Liminal Custody

<https://www.linkedin.com/in/hilallone/>

Q1. Shadow AI used to mean staff pasting data into chatbots. Now it means autonomous agents acting on your systems with no owner. How should a board respond to that shift, and how urgently?

Ans: For Board this means business critical systems become vulnerable to headless operations which has a direct impact on day to day business functions. Absence of accountability and ownership puts these systems at highest risk that can prove catastrophic for business if governance is not enforced. The board needs to make it a zero tolerance issue for agents not associated with a business or human identity. Visibility, detection and inventory becomes critical for boards to track.

Q2. Agent adoption is running far ahead of governance, with only one firm in five prepared. If a CXO could put just one control in place first, what should it be?

Ans: Visibility that leads to comprehensive discovery and inventory of all operational autonomous agents is single most important control that a CXO should put in place.

CXO VOICE – IN CONVERSATION THIS WEEK

Q3. An agent that can read data is one risk. One that can act, send, pay, or delete, is another. Where should the line sit between what an agent may do alone and what needs a human?

Ans: The basic security controls like segregation of duties and roles is equally important for AI agents as it is for human agents. The agent can act with autonomy when boolean or deterministic decisions need to be made. These decisions have a pre-defined guardrails for workflows that allow its execution. A human will need to be involved when it falls outside of the binary permission and guardrails.

Q4. In India, an unsanctioned agent acting on personal data leaves no record for a DPDP audit. How large is that exposure, and what is the most practical way to close it?

Ans: This is a massive gap as data processors may not have visibility into AI agent operations. As with all privacy initiatives, the privacy assessment is primary requirement for an organization that needs to comply with DPDPA. Once privacy impact study is done with AI agents, the recommendations including data anonymization, data isolation and appropriate data usage must be implemented to contain the risk.

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
- In The Wild
- Perspectives
- The Shadow Agent Playbook
- Metric Of The Week
- Fun Fact
- Candid
- **Book Shelf**
- Research Radar
- RACE Pulse
- RACE Programs

RETINA RECOMMENDS · THIS WEEK'S THEME

The Checklist Manifesto: How to Get Things Right


Atul Gawande · Metropolitan Books, 2009

Gawande, a surgeon, set out to answer a simple question. In fields where capable people do complex, high-stakes work at speed, what actually prevents disaster. His answer was not more talent or more training. It was the humble checklist, a short, enforced pause that makes sure the critical steps happen before an irreversible action is taken. Read against this week's cover, the parallel is exact. An agent that can act is a capable worker moving fast, and the thing that keeps it safe is a checkpoint before it does something it cannot take back.

What the book gives a leadership team is permission to be simple. Faced with powerful, fast, autonomous systems, the instinct is an elaborate governance

apparatus. Gawande's evidence, drawn from operating theatres and aircraft cockpits, is that a short list of forcing checks, applied without exception, outperforms grand frameworks that are too heavy to follow. For agents, the reading-versus-acting line in this issue's playbook is exactly that kind of check.

It is not a book about AI, which is its value. It hands you a durable principle, that the discipline is in the pause, and lets you apply it to the newest fast-moving actor in your building. Read it before you design how your agents are allowed to act.



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



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

EDITOR'S NOTE

Two new RACE papers. One puts an adaptive guard in front of your web apps. One lets non-technical staff query enterprise data in plain English, with the trail to prove it. Standard format: a concise technology summary followed by clear enterprise relevance.

Machine Learning-Based Web Application Firewall for Real-Time Threat Detection

Anil Kumar; J. B. Simha; Rashmi Agarwal

Rule-based firewalls only catch attacks they already have a signature for, which leaves a gap every time an attacker varies the pattern. This work builds a web application firewall on an XGBoost model that learns from traffic instead of static rules, detecting SQL injection, cross-site scripting, and local file inclusion in real time. Sitting in front of the web server, it inspects every incoming request, drops the malicious traffic, and passes the rest through, adapting to new attack shapes as it sees them.

Enterprise relevance: It is a working example of the shift this issue argues for, from fixed rules to adaptive, behaviour-based control. For a security team, an ML firewall that learns raises the floor against the fast-changing attacks that static defences miss, and the front-of-server design means it slots into an existing stack without re-architecting.

2024 IEEE Conference on Engineering Informatics (ICEI), Melbourne · IEEE (added to IEEE Xplore, March 2025)

<https://ieeexplore.ieee.org/document/10912239>

LLM-Powered Text-to-SQL Generation and Data Visualization

Ganesh Rohan N; Ratnakar Pandey; Shinu Abhi

Writing SQL against a complex data warehouse is a bottleneck, and most automated text-to-SQL tools produce an answer with no way to trust it. This work tackles both at once. A Graph-based Retrieval Augmented Generation design drives semantically accurate SQL, a generative module shows the query's lineage so a user can see how the answer was built, and Pandas AI handles ad-hoc analysis. In testing it reached up to 91 percent query correctness and cut generation from hours to seconds, letting non-technical users pull from enterprise data and see the working behind each result.

Enterprise relevance: This is how AI widens data access without widening risk. The lineage view is the key move: it turns a black-box answer into one a business user, and an auditor, can trust, which is exactly the explainability that governance now demands. For any firm trying to democratise analytics, it shows how to pair self-serve speed with a trail you can defend.

2026 International Conference on Artificial Intelligence and Data Engineering (AIDE), Nitte · IEEE

<https://ieeexplore.ieee.org/document/11544926>

The CXO takeaway: both papers point the same way, pair AI's speed with a control you can trust, an adaptive firewall that learns faster than attackers, and self-serve data access with a lineage trail an auditor would accept.

Explore the full RACE research archive at <https://raceresearch.reva.edu.in>. For collaboration, licensing or research-partnership enquiries: race@reva.edu.in

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

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

RACEx360 Season 2: that's a wrap

25 September 2026 · Taj West End, Bengaluru ·
with The Economic Times

Season 2 of RACE's Emerging Technology Conference brought 600+ or more builders, buyers and decision-makers across four frontiers: GenAI and Agentic AI, Cloud and Platform Engineering, Cybersecurity in the AI Era, and Deep Tech and Sustainability. It was a day of proof, not slides, with a live Innovation Experience Zone, an AI Hackathon, the EC-Council HACKai Challenge finale, and keynotes that closed with a 90-day action plan. The evening crowned the Top 10 Women Tech Leaders India Awards 2026. Thank you to everyone who joined, spoke, and competed. Highlights and the post-event photos here.

<https://race.reva.edu.in/racex360-season-2>



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

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

OCTOBER 2026 BATCH - APPLY BY 15 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/UGC approved. Industry mentors from Amazon, Microsoft, AWS, EC-Council, Terralogic, IBM and more.

To know more, call:
+91 89040 58866

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

Certified Agentic AI Engineer

31 OCTOBER 2026

Practical expertise in LLMs, AI agents, multi-agent systems, agentic workflows, governance and capstone-based implementation. Real-world use cases, deployable solutions.

To know more, call:
+91 89040 58866

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

RACE SANDBOX SATURDAY 17 OCTOBER 2026

Run your career simulation, live on campus. Sit in a live Master's class, walk the AI, Cybersecurity and Cloud labs, watch participants and alumni demonstrate real projects, and take a 1:1 on where your career goes next. Experience today, decide your tomorrow.

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

Reserve your slot:
+91 89040 58866 ·

enquiry@race.reva.edu.in ·
<https://race.reva.edu.in/>

- From The Editor’s Desk
- The World In 5
- Signal
- Hyperfocal
- In The Wild
- Perspectives
- The Shadow Agent 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

“How many AI agents are already acting inside our systems, who owns each one, and what can they do on their own that we would struggle to undo?”

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