

RACE WEEKLY INTELLIGENCE BRIEF


RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER:

NOBODY ON THE CALL WAS REAL

A finance worker joined a video call with his CFO and colleagues, and wired away \$25 million. Every face on that call was an AI fake, but his. Three seconds of audio is now enough to clone a voice. Seeing and hearing stopped being proof, and your company still runs on both.



ALSO INSIDE

The World in 5 · Three Forces · The Verification Playbook · Metric of the Week · The Trust Signal · Fun Fact · CANDID · Book Shelf · Research Radar · RACE Pulse

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

FROM THE EDITOR'S DESK

In early 2024, a finance employee at the engineering firm Arup joined a routine video call. His CFO was on it, so were several colleagues he recognised. Over the call, he was asked to process a confidential transaction, and he did: fifteen transfers totalling about \$25.6 million. Every other person on that call, including the CFO, was an AI-generated fake. The only real human in the meeting was the one who sent the money.

That case is no longer an outlier; it is the template. The World Economic Forum's 2026 cybersecurity outlook records a quiet but seismic shift: cyber-enabled fraud has overtaken ransomware as the number-one concern of CEOs. Gartner found that 62% of organisations were hit by at least one deepfake attack in the past year, and 37% of security leaders have met a deepfake on a live video call. The reason it is spreading is the collapse in cost: three seconds of audio now produces an 85% voice-clone match, and the tools are free, require no skill, and leave no name.

India is not on the sidelines of this; it is closer to the front line. Indians lost about ₹22,495 crore to cyber fraud in 2025 across 2.8 million complaints; the RBI has reported an eightfold jump in fraud losses driven by AI-generated synthetic identities; and in late July the Supreme Court urged Parliament to write a dedicated criminal law for deepfake-enabled fraud, with a draft bill now being prepared. The IT Rules amended this year already force platforms to pull flagged synthetic content within three hours.

The instinctive response- get better at spotting fakes, or buy a tool that does, is a losing race, and this issue explains why. When humans catch high-quality video deepfakes about a quarter of the time and detection tools lose half their accuracy outside the lab, the answer is not sharper eyes. It is a different discipline: verifying what is real rather than detecting what is fake. Here is how a leadership team rebuilds trust on purpose, before a call that looks exactly like this one costs it everything.

THE WORLD IN 5

GLOBAL HEADLINES A CXO SHOULD BE ABLE TO REFERENCE

- From the Editor's Desk
- **The World in 5**
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- 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 CASE:

An entire call, faked

The Arup incident: roughly \$25.6 million transferred across fifteen transactions after a video call on which every participant except the victim was AI-generated; remains the defining example of deepfake fraud, and it has been widely followed since. Ferrari narrowly avoided a similar attack when an executive, called by a voice cloned to match the CEO's, asked a personal question the impersonator could not answer. The difference between a headline loss and a near-miss was one verification question.

2. THE SHIFT:

Fraud passes ransomware

The World Economic Forum's Global Cybersecurity Outlook 2026 found that cyber-enabled fraud has overtaken ransomware as the top concern for chief executives. Gartner's data explains why: 62% of organisations reported at least one deepfake attack in the past twelve months, and 37% of security leaders have personally encountered a deepfake on a live video call. The threat has moved from the security team's dashboard to the CEO's own inbox and phone.

3. THE EASE:

Zero cost, zero skill

Three seconds of audio is now enough to generate a voice clone with about 85% accuracy; and most executives have hours of public audio in conference recordings, interviews and webinars. The 2026 International AI Safety Report noted that the tools powering these scams are free, need no technical expertise, and can be used anonymously. That combination- no cost, no skill, no accountability- is why this threat is growing faster than any other.

4. INDIA:

A national fraud epidemic

Indians lost roughly ₹22,495 crore (about \$2.7 billion) to cyber fraud in 2025 across 2.8 million complaints, up 24% year on year. The RBI has flagged an eightfold rise in fraud losses driven by AI-generated synthetic identities, with real-time face-swaps used to spoof video-KYC and cloned voices to bypass OTPs. Deepfaked celebrity endorsements- Tendulkar, Tata, Kohli and others- now routinely front investment scams. In late July, the Supreme Court pressed Parliament to criminalise deepfake-enabled fraud directly.

5. THE DETECTION TRAP:

The tools are losing

Humans identify high-quality video deepfakes only about a quarter of the time, and one large study found just 0.1% of people could reliably distinguish every real and fake sample shown to them. Detection tools do better in the lab – up to 96%, but lose 45 to 50% of that accuracy in real-world conditions. The uncomfortable conclusion the industry is reaching: you cannot detect your way out of this, and the defence has to shift from spotting fakes to proving what is real.

Source:

<https://www.weforum.org/publications/global-cybersecurity-outlook-2026/>

<https://www.gartner.com/en/newsroom/press-releases/2025-09-02-why-cios-cannot-ignore-the-rising-tide-of-deepfake-attacks>

https://www.mcafee.com/en-us/newsroom/press-releases/press-release.html?news_id=5aa0d525-c1ae-4b1e-a1b7-dc499410c5a1

https://145615497.fs1.hubspotusercontent-eu1.net/hubfs/145615497/Gated%20PDFs/iProov_Threat%20Intelligence%20report_1%202025%20-.pdf

<https://rbi.org.in/> <https://www.sci.gov.in/>

 SIGNAL

THE AUTHENTICITY COLLAPSE — BY THE NUMBERS

- From the Editor's Desk
- The World in 5
- **Signal**
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

\$25.6M

transferred by one Arup employee after a video call on which every other participant was an AI fake

3 sec

of audio now enough to clone a voice at ~85% accuracy — and most executives have hours of it online

62%


of organisations hit by at least one deepfake attack in the past year; 37% on a live video call (Gartner)

24.5%

the rate at which humans correctly spot a high-quality video deepfake — worse than a coin toss

₹22,495 cr

lost to cyber fraud in India in 2025 across 2.8 million complaints, up 24% year on year



HYPERFOCAL

WHAT BROKE THIS WEEK

- From the Editor's Desk
- The World in 5
- Signal
- **Hyperfocal**
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

You Cannot Detect Your Way Out of This

The natural response to fake faces and cloned voices is to get better at spotting them, or to buy a tool that does. Both are losing bets, and the numbers say so plainly. Humans catch high-quality video deepfakes about a quarter of the time, worse than a coin toss, and the single best study found only one person in a thousand could reliably tell every real sample from every fake one. Detection software looks impressive in the lab, but loses roughly half its accuracy in the messy conditions of a real call, and every month the generators improve faster than the detectors.

This is an arms race the defender structurally loses, because the same models power both sides and the attacker only has to win once. Spending the security budget on sharper detection is like buying a faster horse in a week the tractor arrived: it answers the

wrong question. The right question is not “how do we spot the fake?” but “how do we prove what is real?”, and that is a different discipline entirely, built on verification and provenance rather than perception.

The shift is from detection to authentication. Instead of asking a human or a tool to judge whether a face or voice is genuine- a judgment both now fail- you verify identity and authority through a channel the attacker does not control: a call-back to a known number, a shared secret, a second approver, a cryptographic signature on a document. The Arup loss and the Ferrari near-miss differed by exactly this. One trusted the call; the other asked a question the fake could not answer. Trust stopped being something you can perceive. It has to be something you verify.

RETINA TAKE

For a century, seeing and hearing were proof. That assumption is now false, and every payment approval, identity check and executive instruction that still rests on it is an open door. The organisations that come through this will be the ones that rebuilt trust as a verified process, before a call that looked completely real asked them for money.

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- **Panorama**
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

PANORAMA

FORCE 1

The Cost of a Convincing Fake Has Fallen to Zero

What made the Arup-style attack rare was effort; that barrier is gone. Three seconds of audio produces a usable voice clone, real-time face-swaps run on a live call, and the tools are free, skill-free and anonymous. Every executive with a public profile, a keynote, a podcast, an earnings call, has already supplied the raw material to be cloned. This is no longer a threat aimed only at the largest firms; it is a commodity attack that scales to anyone with a payment approver and a CEO whose voice exists online.

The Move: Assume your executives are already cloneable and plan from there. The defensive question is not whether someone can fake your CFO, they can, but whether a faked CFO can actually move money or data through your process. Close that path, because the fake is now free and inevitable.

FORCE 2

Detection Is a Race the Defender Loses

Humans and tools both fail at spotting fakes in the wild, and the generators improve faster than the detectors, so any strategy that depends on catching the fake is depreciating by the month. The winning move is to stop judging authenticity by perception and start proving it by verification, identity and authority confirmed through a channel the attacker cannot forge. This is the same zero-trust logic the security world already applies to networks, now applied to people: verify explicitly, trust nothing

THREE FORCES — AND THE MOVE EACH ONE DEMANDS

by default, not even a familiar face on a call.

The Move: Move the budget and the training from detection to verification. Teach the organisation that the defence is not a sharper eye but a fixed habit: high-stakes instructions are confirmed out-of-band, every time, regardless of how real the request looks or who appears to be making it.

FORCE 3

When Anything Can Be Faked, Everything Real Becomes Deniable

There is a second-order cost beyond the direct fraud, and it is arguably larger. As fakes become undetectable, genuine content loses its authority too, a real recording can be dismissed as a deepfake, and a real instruction doubted. Analysts call it the liar's dividend, or truth decay: once everything can be faked, nothing can be trusted at face value, and every video call, voice message and email becomes suspect. For an enterprise, that erodes the speed and trust its operations quietly depend on, from approvals to customer communication to brand.

The Move: Invest in provenance for your own communications, verified channels, content credentials, signed and authenticated outbound messages, so your customers and staff can confirm what genuinely came from you. In a world of cheap fakes, being provably real is a competitive advantage, not just a control.



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- **Perspectives**
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

Stop Trying to Spot the Fake



**Sandeep
Vijayaraghavan**

EVP, Cybersecurity and
Cloud Services, Terralogic

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

When Seeing Is No Longer Believing: Deepfakes and the New Face of Financial Fraud

For years, we have been told not to trust every email, click on unknown links, or share our OTPs. But artificial intelligence is forcing us to rethink one of our most basic assumptions: **can we trust what we see and hear?** Deepfake technology has enabled the creation of remarkably convincing videos and voices of real people. What concerns me most as a cybersecurity professional is not the technology itself, but how quickly it is being weaponised to exploit something far

more powerful than a computer vulnerability: **human trust.**

India has already seen a disturbing example. In Pune, a 59-year-old man reportedly came across an Instagram advertisement featuring deepfake videos of Infosys co-founder **N. R. Narayana Murthy and Sudha Murty**, apparently promoting an AI-based share-trading platform. The advertisement promised extraordinary returns from a small initial investment. Believing the videos to be genuine, the victim initially invested ₹21,000. Over the following months, fraudsters repeatedly contacted him, persuaded him to make additional transfers and even encouraged him to borrow money. According to the police complaint, he eventually made 19 transfers totalling about **₹43.68 lakh** before realising that the entire investment platform was fraudulent.

What makes this case particularly troubling is that the fraud did not depend on a sophisticated piece of malware. It depended

on **credibility.** The attackers used familiar faces, convincing voices, social media advertising, fake investment information and persistent human interaction to build confidence. The technology effectively became a digital actor playing a trusted person. Fact-checkers have confirmed that videos using the Murty's likenesses were manipulated, including AI-generated audio, and that the original source material did not contain the investment claims shown in the fraudulent versions.

This is where I believe we need to change the way we think about cybersecurity. Traditionally, we have focused on protecting devices, networks, applications and passwords. Deepfake fraud attacks a different layer: **our judgement.** A fraudster can now combine a cloned voice, manipulated video, fake social-media profile, fraudulent website and a sense of urgency to create an entire deception ecosystem. By the time the victim sees the red flags, the emotional decision may already have been



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- **Perspectives**
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

made. Even Sudha Murty has subsequently warned publicly that fake videos using her image and voice were being used to promote financial schemes and urged people not to make investment decisions based on such content.

So what should an ordinary user do? My advice is simple: **pause, question and independently verify**. Never treat a video or voice recording as proof of identity. Be especially suspicious of guaranteed or unusually high returns, urgent investment

opportunities and requests to move conversations to WhatsApp or Telegram. Check whether the information appears on the person's verified or official channels. Most importantly, before transferring money, verify the investment platform, beneficiary and claims through an independent source, not through a link or phone number provided by the person who contacted you.

Deepfakes are not going away; they will become better, cheaper and harder to detect. The answer therefore cannot be to expect

every citizen to become an AI expert. Banks, financial institutions, technology companies and governments need stronger layers of **identity verification, AI-generated-content detection, behavioural analytics and transaction monitoring**. But technology alone will not solve this problem. As we enter an era where seeing is no longer believing, our most important cybersecurity control may be a simple human habit: **stop, verify and then act**. Our money should never move simply because a familiar face told us to move it.

RACE PERSPECTIVE

The deepfake era does not reward the sharpest eye; it rewards the firmest process. You will not out-detect the fakes, no one will. You will out-verify them, by making "confirm through a channel they cannot control" the default for anything that moves money, data or trust.



- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- **The Verification Playbook**
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE VERIFICATION PLAYBOOK

REBUILDING TRUST AS A PROCESS,
NOT A PERCEPTION

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

MOVE 1

MAKE OUT-OF-BAND VERIFICATION MANDATORY FOR MONEY AND AUTHORITY

The single highest-return control: no high-value payment, no change to banking details, and no sensitive-data release acts on a voice or video instruction alone. It is confirmed through an independent channel, a call-back to a pre-verified number, a second approver on a separate system. The Ferrari attack failed on exactly this; the Arup attack succeeded for want of it.

The Move: Write the rule this week and make it unskippable: high-stakes instructions require out-of-band confirmation, always, with no exception for seniority or urgency. The rule protects the approver as much as the company.

MOVE 2

KILL “URGENCY PLUS SECRECY” AS A VALID INSTRUCTION

Nearly every deepfake fraud runs on the same script: this is urgent, it is confidential, do not loop anyone in, act now. That combination is not a sign of importance; it is the signature of the attack, engineered to rush the target past the verification step. Teach the whole organisation to treat urgency-plus-secrecy as a trigger to slow down and verify, not to comply.

The Move: Train staff to recognise the pattern, not the fake. “Urgent, secret, act now” from any executive is the cue to stop and confirm through a second channel; the more pressure to skip verification, the more it is needed.

MOVE 3

HARDEN IDENTITY ONBOARDING AGAINST SYNTHETIC FACES AND VOICES

Real-time face-swaps now spoof video-KYC, and cloned voices defeat voice-OTP, so verification steps built for humans are being beaten by machines. Where you verify identity, customer onboarding, account recovery, high-value authorisation, assume the face and voice may be synthetic, and add liveness and provenance checks designed for that reality. This is where India's synthetic-identity fraud is concentrated.

The Move: Audit every identity checkpoint for its resilience to a real-time deepfake, not just a stolen password. The RBI's synthetic-identity warnings are a map of where to look first.

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- **The Verification Playbook**
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE VERIFICATION PLAYBOOK

REBUILDING TRUST AS A PROCESS,
NOT A PERCEPTION

MOVE 4

GIVE YOUR OWN COMMUNICATIONS PROVENANCE

Defence is not only about incoming fakes; it is about letting people confirm what genuinely came from you. Verified channels, content credentials on official media, signed communications- these let customers and staff distinguish the real you from a clone. As deepfaked executive and brand impersonation spreads, being provably authentic protects both revenue and reputation.

The Move: Establish and publicise the channels through which you will, and will never, contact customers and staff, and adopt content-provenance standards for official media. Make “how to check this is really us” easy to find.

MOVE 5

Rehearse the incident and know the golden hour

When a fake succeeds, speed of response decides how much is recovered. In India the first hour is critical, reporting immediately to 1930 or cybercrime.gov.in can let banks freeze the receiving account before the money is layered away. A verification failure should trigger a rehearsed response, not improvisation, and staff should know the number before they need it.

The Move: Add a deepfake-fraud scenario to your incident-response drill, and put the reporting steps and the golden-hour rule where finance and frontline staff will find them fast. Practised recovery is the difference between a scare and a loss.

Source:

<https://www.weforum.org/publications/global-cybersecurity-outlook-2026/digest/>

<https://www.gartner.com/en/newsroom/press-releases/2025-09-02-why-cios-cannot-ignore-the-rising-tide-of-deepfake-attacks>

<https://www.rbi.org.in/> <https://www.cybercrime.gov.in/>

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- **Metric Of The Week**
- **Fun Fact**
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

METRIC OF THE WEEK

ONE NUMBER, PROPERLY UNDERSTOOD

3 seconds

of audio is now enough to clone a voice

What it is: McAfee's finding that roughly three seconds of recorded speech can generate a voice clone with about 85% accuracy, down from the hours of source audio needed only a few years ago. The barrier that once made voice impersonation rare has effectively disappeared.

Why it matters now: Every executive with a public footprint, a conference talk, a podcast, an earnings call, a voicemail


greeting, has already provided far more than three seconds. The raw material for cloning your leadership is not something an attacker must obtain; it is already published. The clone is not a risk you can prevent, only one you can plan around.

How leaders use it: Assume every senior voice in your organisation is already cloneable, and test the consequence: if a perfect clone of your CEO called your

finance team right now with an urgent request, what would stop the money leaving? The answer should be a process, not a person's ear.

Source:

https://www.mcafee.com/en-us/newsroom/press-releases/press-release.html?news_id=5aa0d525-c1ae-4b1e-a1b7-dc499410c5a1t/



FUN FACT

In 1950, Alan Turing proposed a test for machine intelligence: could a computer, through conversation alone, convince a human it was another human? For seventy-five years, passing the "imitation game" was treated as a milestone; the moment a machine could be said to think. It has now been passed, comprehensively, and the achievement did not arrive as a philosophical triumph. It arrived as a fraud technique. The exact capability Turing named as the definition of intelligence, a machine impersonating a person well enough to fool one, is the thing draining tens of thousands of crores a year. The test became the weapon.

Source: Alan Turing, "Computing Machinery and Intelligence," 1950

SOMETHING TO REPEAT AT DINNER

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- **The Trust Signal**
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

⚡ THE TRUST SIGNAL

INDIA'S AUTHENTICITY NUMBERS

The fraud and regulation data that will shape your risk posture whether you track it or not.

₹22,495 cr

lost to cyber fraud in India in 2025 (about \$2.7 billion), across 2.8 million complaints — up 24% year on year.

8x

the reported rise in bank fraud losses driven by AI-generated synthetic identities (RBI).

3 hours

the window platforms have to remove flagged synthetic content under India's IT Rules amendment, in force since February 2026.

Jul 2026

the Supreme Court urged Parliament to enact a dedicated criminal law for deepfake-enabled fraud; a draft bill is being prepared.

1930

India's cyber-fraud helpline — reporting inside the "golden hour" lets banks freeze the receiving account before funds are moved.

0.1%

of people, in one large study, could reliably tell every real sample from every fake one they were shown.

Source: <https://www.mha.gov.in/> · <https://i4c.mha.gov.in/> · <https://www.rbi.org.in/> · <https://www.meity.gov.in/> · <https://www.sci.gov.in/>
https://145615497.fs1.hubspotusercontent-eu1.net/hubfs/145615497/Gated%20PDFs/iProov_Threat%20Intelligence%20report_1%202025%20-.pdf

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- **Candid**
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



CANDID



Dr. Ram Kumar G Ph.D, CISM, PMP
Cyber Security & Risk Leader,
Global Automotive Company

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

Q1 A finance worker wired \$25 million because everyone on the video call looked and sounded real. Have most organisations accepted yet that a familiar face on a call is no longer proof of anything, or is that still sinking in?

Ans: Most organizations are still playing catch-up, with deepfake technology evolving far faster than corporate awareness. While security experts have long warned that synthetic media invalidates visual trust, the majority of companies still rely on outdated visual and voice recognition instead of out-of-band verification.

Until cryptographic identity checks and multi-channel protocols become standard, high-stakes fraud like the \$25 million deepfake heist will continue to catch businesses off guard.

Q2 Detection tools lose half their accuracy in the real world. Is the answer better detection, or verifying identity through channels an attacker can't fake – and what does that shift look like in practice?

Ans: The solution is to shift away from detection toward mandatory, out-of-band verification through un-fakable channels. Because AI detection tools struggle in real-world conditions and is not yet foolproof, security relies on out-of-band protocols—such as secondary confirmation via encrypted messaging or hardware keys.

In my view, true resilience requires treating visual and vocal familiarity as untrusted data until verified through strict cryptographic identity checks.

Remi, in today's AI era, with deepfakes mimick reality, what you see isn't what you get!

Q3 India loses tens of thousands of crores a year to this, and the Supreme Court is pushing for a criminal law. What can enterprises do that regulation and law cannot?

Ans: Indian citizens have lost a whopping Rs 10,000+ crores to various cyber crimes in H1 2026, as per report appearing in the Indian Express.

Enterprises can implement proactive defenses that function faster than legislative enforcement. While laws punish after the cyber crime happens and the criminals are convicted, businesses can mandate real-time cryptographic verification, hardware token authentication, and multi-person authorization protocols for financial transfers. Business resilience relies on zero-trust architectures and continuous employee awareness to stop fraud at the point of attempt.

Q4 Final thoughts on doing business in a world where you can no longer trust the voice on the phone or the face on the screen.

Ans: Doing business today requires replacing default visual trust with cryptographic verification and zero-trust processes. Survival in a deepfake-driven world means treating identity as something to be authenticated, not assumed.

I'd say the old saying 'Seeing is believing' no longer holds good in the era of deepfakes. Ultimately, organizational resilience lies not in fearing synthetic media, but in building workflows where no single voice or video call has unchecked authority.

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
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- **Book Shelf**
- Research Radar
- RACE Pulse
- RACE Programs

RETINA RECOMMENDS · THIS WEEK'S THEME

Deepfakes: The Coming Infocalypse

Nina Schick · Twelve, 2020

Schick, an advisor to governments and technology leaders on synthetic media, wrote the warning six years before the losses arrived. Her argument was that AI-generated media would not merely enable new scams but corrode the shared reality that institutions, markets and democracies rest on — an “infocalypse” in which the very idea of authentic evidence comes under question. In 2020 it read as a provocative forecast. In 2026 it reads as a description of the morning news.

Read now, with the Arup case and India's fraud numbers in hand, the book's value is that it gets ahead of the individual scam and onto the systemic problem — the liar's dividend, the erosion of trust in all recorded media, the strategic use of fakes by states

and criminals alike. It is the conceptual map behind this week's cover: not “how do I spot one deepfake” but “what happens to an organisation, and a society, when nothing recorded can be trusted by default.”

It is a briefing, not a manual, and deliberately short. What it gives a leadership team is the vocabulary and the frame to take this seriously as a strategic risk rather than an IT nuisance — which is the first step toward building the verification discipline the rest of this issue describes. Read it before your next board conversation about fraud.



- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- **Research Radar**
- RACE Pulse
- RACE Programs



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

EDITOR'S NOTE

Two new Springer publications from RACE — one bringing audio analytics to India's regional music, one strengthening the machine-learning defences behind cloud security. Standard format: a concise technology summary followed by clear enterprise relevance.

Audio Analytics · Machine Learning · Recommender Systems
Kannada Music Genre Classification Using Audio Analytics, Machine Learning and Deep Learning

B. S. Mallikarjuna; Jay B. Simha; Shinu Abhi

Streaming recommenders still lean on metadata — album, film, singer, year — and miss the finer genres that define regional traditions. This work builds, apparently for the first time, a Kannada soft-music dataset of 1,500 audio clips and applies audio-feature extraction with ML and deep-learning models to classify genres such as Janapada, Bhavageethe and Classical, reaching 96.05% accuracy after PCA, cross-validation and hyperparameter tuning. Proposed uses include a genre-aware recommender and a "teacher-in-machine."

Enterprise Relevance: For media, streaming and edtech players in India's language markets, this shows how audio analytics can unlock regional catalogues that metadata alone leaves invisible — better discovery, deeper engagement, and a route into audiences generic engines under-serve.

ICDSA 2024 · Data Science and Applications · Lecture Notes in Networks and Systems, vol 1263, pp 521–536 · Springer, Singapore, 2025
https://link.springer.com/chapter/10.1007/978-981-96-2724-0_38

Cloud Security · Anomaly Detection · Machine Learning
Enhanced Anomaly Detection in Secure Cloud Networks Using Machine Learning

Shameer Babu; Pradyumn Nand; Shinu Abhi

As cloud threats multiply and signatures lag, this work tests machine learning as a more adaptive line of defence. Using the UNSW-NB15 intrusion dataset, it engineers diverse feature sets, generates six randomised training sets simulating varied attack scenarios, and benchmarks six anomaly-detection algorithms on accuracy, precision, recall and F1 — with Random Forest leading consistently at about 97.82%.

Enterprise Relevance: For CISOs and cloud teams, it reinforces this week's cover thesis from the systems side: when static rules can't keep up, behaviour-based detection catches what signatures miss — a pattern that applies to fraud and identity anomalies as much as to network intrusions.

ICDSA 2024 · Data Science and Applications · Lecture Notes in Networks and Systems, vol 1266, pp 105–121 · Springer, Singapore, 2025
https://link.springer.com/chapter/10.1007/978-981-96-2647-2_7

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

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

RACEx360 SEASON 2

25 SEPTEMBER 2026 · TAJ WEST END, BENGALURU

One day, one stage, four frontiers.

REVA Academy's Emerging Technology Conference returns, in partnership with The Economic Times — 750+ builders, buyers and decision-makers across four tracks that matter: GenAI & Agentic AI · Cloud & Platform Engineering · Cybersecurity in the AI Era · Deep Tech & Sustainability.

Not a day of slides — a day of proof: an Innovation Experience Zone with live demos across AI, robotics, cybersecurity, cloud and quantum; live build sessions; an AI Hackathon and Cybersecurity Challenge; and keynotes that close with a 90-day action plan you can take back to your team.

The evening turns gold with the Top 10 Women Tech Leaders India Awards 2026, now in its 4th edition — India's only annual awards programme dedicated to women technology leaders.

Early Bird 4,000 + GST until 15 August · Standard 5,000 + GST. Passes cover all four tracks, the Expo Hall and the Experience Zone.

GET YOUR PASS:

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

BECOME A SPONSOR:

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



RACEx360
Emerging Technology Conference 2026

In partnership with  **THE ECONOMIC TIMES**

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

REVA Academy for Corporate Excellence

THE ENTERPRISE

GAVE AI THE KEYS

You get 90 minutes to prove it shouldn't have.

EC-Council

HACKai

Challenge

Prompt injection. RAG poisoning. Jailbreaking.
Agentic privilege escalation.

QUALIFIER

10th to 18th September 2026

FINALE

25th September 2026 at
The TAJ West End, Bengaluru

REVA
UNIVERSITY
Bengaluru, India

CASH PRIZE

₹1.75 Lakhs

+
CEH Vouchers

REGISTER TODAY

Registration closes on 10th September 2026

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

SEPTEMBER 2026 BATCH · APPLY BY 10 AUGUST 2026

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.

To know more, call:
+91 89040 58866

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

NEW BATCH **Certified Agentic AI Engineer** 1 AUGUST 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 5 SEPTEMBER 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
- Panorama
- Perspectives
- The Verification Playbook
- Metric Of The Week
- Fun Fact
- The Trust Signal
- 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

"If a perfect deepfake of our CEO called finance right now with an urgent transfer, what would actually stop the money leaving — a process, or someone's ear?"

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