

RACE WEEKLY INTELLIGENCE BRIEF


RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER:

YOU DON'T OWN WHAT YOUR AI MADE

Purely AI-generated work has no copyright owner, so your competitor can copy it freely. But if it infringes someone else's work, you can still be sued. You get the liability without the ownership, and almost everything your company shipped with AI this year sits in that gap.



ALSO INSIDE

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

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

FROM THE EDITOR'S DESK

Last week the machine attacked your systems. This week it comes for something quieter and, for most enterprises, less understood: the ownership of everything it makes for you.

Two landmark rulings landed within four days of each other, on opposite sides of the world. On 20 July, a US federal judge gave final approval to a \$1.5 billion settlement, the largest copyright class action in American history, over the books used to train an AI model. On 24 July, the Delhi High Court, in a 135-page order after 32 hearings, refused a news agency's bid to stop OpenAI training on its archive, holding on a preliminary basis that such training can be fair dealing under Indian law. Read together, and alongside the US Supreme Court's refusal in March to revisit the rule that only humans can hold copyright, they mark the moment the law stopped being theoretical about AI and intellectual property.

The picture they paint is not the one most boardrooms assume. The comfortable assumption is that AI output is a productivity gift you own and can defend. The law now says something stranger and more uncomfortable. Work generated purely by AI, with no meaningful human authorship, has no copyright owner at all, which means your competitor can lift your AI-written product copy, your AI-generated designs, even AI-authored code, and you cannot stop them. And yet if that same output happens to reproduce someone else's protected work, the liability can still land on you. You hold the risk of the thing you do not own.

Meanwhile the safety net most enterprises are counting on, vendor copyright indemnities, turns out to be stitched with conditions few teams actually meet, and the insurance behind it is being quietly withdrawn. This issue maps the ownership trap, the liability you cannot see, and the indemnity you think you have — and gives leadership five moves to take before the asset you built on AI turns out to be one you can neither protect nor defend.

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 IP Playbook
- Metric Of The Week
- Fun Fact
- The IP 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. **OWNERSHIP: only humans can hold a copyright**

On 2 March 2026, the US Supreme Court declined to hear Thaler v. Perlmutter, leaving intact the rule that a work must have a human author to be copyrightable. The practical consequence for enterprises is blunt: output generated purely by AI, without meaningful human creative control, is not eligible for copyright, it has no owner, and anyone can copy and use it freely. AI-assisted work can still be protected, but only where a human shaped the expressive result; a detailed prompt, on its own, is not enough.

2. **TRAINING: the first real price on the data**

On 20 July 2026, a US federal judge gave final approval to Anthropic's \$1.5 billion settlement with authors and publishers, about \$3,000 for each of roughly 500,000 books, over pirated copies used in training. A 2025 ruling in the case had found that training on lawfully acquired books was fair use, but that storing pirated ones was not. Crucially, because the case settled rather than went to appeal, it sets no binding precedent — it sets a price. Parallel suits against OpenAI, Google, Meta, Getty and others remain live, including a fresh

class action filed against Google over Gemini on 10 July.

3. **INDIA: fair dealing, for now, but jurisdiction lands**

On 24 July 2026, the Delhi High Court refused ANI Media's interim injunction against OpenAI, holding on a prima facie basis that training ChatGPT on ANI's news archive can fall within India's fair-dealing exception. But the court also rejected OpenAI's core defence: it held that Indian courts do have jurisdiction even though the training happened on foreign servers, reasoning that a platform serving Indian users cannot escape Indian law through server location. The findings are provisional and the main suit continues, but the burden now sits with rights-holders to prove reproduction, not merely ingestion.

4. **THE DOCKET: output infringement is still open**

Every ruling so far addresses the input side, whether training on copyrighted material is lawful. The output side, whether what an AI generates for you can infringe someone else's work, and who is liable when it does, remains largely unresolved across the active US cases

(NYT v OpenAI, Getty v Stability, the music suits against Suno and Udio). This is the exposure that reaches ordinary enterprises, not just AI labs: you are the one publishing the output.

5. **THE SAFETY NET: indemnities narrow, insurers exit**

The vendor copyright indemnities enterprises rely on Microsoft's Copilot Copyright Commitment, OpenAI's Copyright Shield and their peers- apply only to paid enterprise tiers and are voided if you disable content filters, tamper with safety systems, or use output you had reason to know was infringing; most exclude trademark entirely. At the same time, insurers have moved the other way: standard policy endorsements taking effect in January 2026 introduced broad generative-AI exclusions, and at least one carrier added an 'absolute' AI exclusion to directors-and-officers cover.

Source: <https://supremecourt.gov/> · <https://reuters.com/> · Delhi High Court, ANI Media v OpenAI · <https://copyright.gov/> · vendor terms of service



SIGNAL

THE OWNERSHIP TRAP — BY THE NUMBERS

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

ZERO

copyright protection on purely AI-generated output under current US law, no human author, no owner, freely copyable

\$1.5 bn

the Anthropic training-data settlement, largest copyright class action in US history, ~\$3,000 per book

24 JUL

Delhi High Court's prima facie ANI v OpenAI ruling, fair dealing provisionally, but Indian jurisdiction affirmed

2 MAR

the date the US Supreme Court let the human-authorship rule stand (Thaler v. Perlmutter)

JAN 2026

when broad generative-AI exclusions began entering standard insurance policies, as vendor indemnities narrowed



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

HYPERFOCAL

WHAT BROKE THIS WEEK

The Asset You Can Neither Protect Nor Fully Defend

Put the two rulings side by side and a strange shape appears. The Anthropic settlement and the ANI decision are both about the input, the data that goes into training a model. Neither touches the question that actually reaches an ordinary company: who owns, and who is liable for, the output that comes out. And on that question, the law that does exist points in an awkward direction.

Start with ownership. The human-authorship rule, reaffirmed by the Supreme Court's refusal to disturb it, means that a marketing image, a block of code, or a product description generated purely by prompting an AI has no copyright owner. Not you, not the vendor, not the model, no one. In the words of practitioners reading the 2026 position, it belongs to no one, and anyone can copy and use it freely.

The polished campaign your team generated in an afternoon is, legally, a thing your competitor may reproduce without asking.

Now turn it over. That same output, which you cannot protect, can still expose you. If what the model generated reproduces someone else's protected work closely enough, a real risk, since models are trained on protected material, the infringement claim lands on the party who published it. That is you. So the enterprise sits in a genuinely perverse position: the more of its work product is purely machine-made, the less of it the enterprise owns, and the more of it carries a liability the enterprise cannot see. Ownership and liability, which normally travel together, have been pulled apart.

RETINA TAKE

The instinct is to treat AI output as free productivity. The law treats it as an orphan you are nonetheless responsible for. Until an enterprise knows which of its AI-made work it actually owns, and which it merely published, it is managing an asset and a liability it has never counted.

PANORAMA

THREE FORCES — AND THE MOVE EACH ONE DEMANDS

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

FORCE 1

The Output You Cannot Protect

The more your workflow leans on pure generation, prompt in, deliverable out, the more of your work product falls outside copyright entirely. Purely AI-generated code, copy, art and design have no owner and can be freely copied by anyone, including the competitor you were trying to out-produce. The speed that makes generative AI attractive is the very thing that strips the protection: the less a human shapes the expressive result, the less there is to own. For any business whose edge is its creative or written output, this is not a footnote; it is a hole in the asset base.

The Move: Decide which AI-assisted work needs to be a protected asset and which is disposable. For the work that matters- the brand campaign, the flagship product's code, the signature design- keep a human demonstrably in the creative loop and keep the record that proves it. Protection now depends on evidence of human authorship, not on the tool used.

FORCE 2

The Liability You Cannot See

Every AI-copyright ruling to date is about training. The question of whether a specific output infringes someone else's work, and who pays when it does, is still open, and it is the one that reaches you, because you are the publisher. A model trained on protected material can reproduce fragments of it; when the enterprise ships that output as its own, the exposure is yours, not the lab's. This is the inverse of Force 1: the work you cannot own, you can still be sued over.

FORCE 3

The Indemnity You Think You Have

Most enterprises believe their vendor contract covers them. It may not. The major copyright commitments apply only to paid enterprise tiers, and they are voided by exactly the things productivity-minded teams do, disabling or loosening content filters, tampering with safety systems, or using output you had reason to suspect was infringing. Trademark is typically excluded outright. As one lawyer put it, the vendors are grading their own homework: the indemnity is only as strong as your ability to prove you met its conditions. And the backstop is thinning, insurers began adding broad generative-AI exclusions to standard policies in 2026, so the gap the indemnity leaves may not be covered elsewhere.

The Move: Have someone read the actual carve-outs in your AI vendor agreements, then make meeting them a policy: keep content filters on by default, document that you did, and confirm the indemnity covers the tier and use case you actually run. Separately, ask your broker whether your current cover still responds to an AI-related IP claim, the answer has quietly changed.

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

PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

DEEP TECH — ANALYSIS BY THE RACE TEAM

By Dr. J. B Simha, Chief Mentor, AI REVA Academy for Corporate Excellence, REVA University, Chief AI Officer, AdoptAI

For decades, technology companies measured their intellectual capital through patents, software, documents, and proprietary algorithms. The assumption was simple: if you owned the code, you owned the advantage. That assumption is becoming less reliable.

Generative AI can produce thousands of lines of working code in minutes. Open-source models continue to narrow the gap with proprietary systems. Software frameworks evolve faster than product lifecycles. In such an environment, code itself is becoming easier to generate and easier to replace. What remains difficult to replicate is something else entirely. It is the enterprise's accumulated intelligence.

That intelligence lives in proprietary datasets, domain expertise, business rules, evaluation methods, feedback from customers, model tuning, AI agents, operational workflows, and the decisions made over years of solving real business problems. These are the assets that determine whether two organizations using the same foundation model produce ordinary

results or industry-leading ones. This changes how we should think about intellectual property.

The debate surrounding AI has largely centered on whether AI-generated content or code can be copyrighted. That question matters, but it is not the first question a leadership team should ask. Copyright protects certain forms of expression. It does not, by itself, protect business insight, organizational know-how, proprietary data, or the performance of a carefully engineered AI system.

For many enterprises, those assets are considerably more valuable than the source code generated along the way. The challenge, therefore, is not simply legal. It is architectural.

Every AI system is built on a sequence of decisions. Which data is collected? How is it cleaned? Which models are approved? How are they evaluated? What human expertise is incorporated? Which outputs are accepted? What evidence is retained? These decisions

shape the quality, reliability, and defensibility of the resulting system.

Architecture determines far more than technical performance. It determines what the enterprise can protect.

Consider two organizations building AI solutions for the same business problem.

Both use the same foundation model. Both generate similar amounts of code. Yet one has spent years curating proprietary datasets, capturing expert knowledge, validating outputs against domain-specific benchmarks, refining prompts, measuring model performance, and continuously improving its AI agents through operational feedback. The other depends almost entirely on publicly available information and generic prompts.

The software may look similar. The enterprise value does not. The competitive advantage belongs to the organization that has accumulated intelligence, not merely generated code. This is where governance becomes a strategic capability rather than a compliance obligation.



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

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

An enterprise that maintains disciplined control over its data assets, model inventory, evaluation pipelines, human review processes, deployment policies, and operational records is not simply reducing regulatory risk. It is preserving institutional knowledge. It is creating evidence of how decisions were made. It is strengthening customer trust. It is protecting assets that competitors cannot easily reproduce.

Good governance therefore contributes to security, quality, resilience, and intellectual property simultaneously.

The legal landscape continues to evolve. Courts across different jurisdictions are still defining the boundaries of AI-assisted authorship, copyright, liability, and training practices. Indian enterprises should expect these questions to become more prominent as AI adoption accelerates. Rather than waiting for legal certainty, organizations

should build processes that demonstrate responsible human oversight, clear provenance, and disciplined AI governance. Those investments retain their value regardless of how individual legal questions are ultimately resolved.

Perhaps the most important shift is conceptual. The valuable output of an AI programme is rarely the code itself. It is the capability the enterprise has created.

A fraud detection model that consistently outperforms competitors. A clinical decision-support system refined through years of specialist feedback. A manufacturing optimization platform trained on proprietary production data. An intelligent customer-support agent that captures decades of organizational knowledge.

These systems derive their value from the intelligence embedded within them, not merely from the software that implements

them.

Technology leaders should therefore ask different questions.

- Which knowledge assets make our AI systems unique?
- Which datasets would our competitors struggle to obtain?
- How do we preserve the expertise that improves our models?
- Can we explain how our AI systems reach their conclusions?
- What mechanisms protect these assets if employees leave, vendors change, or regulations evolve?

Organizations that can answer these questions are building sustainable competitive advantage. Those that cannot may discover that while AI made software easier to create, it also made software alone a less durable source of differentiation.

RACE PERSPECTIVE

The strategic asset in the AI era is not generated code alone. It is the enterprise intelligence embodied in data, knowledge, models, evaluation processes, AI agents, and disciplined workflows. Copyright, patents, trade secrets, contracts, and governance each play a role in protecting that intelligence, but none is sufficient in isolation. The organizations that will lead are those that deliberately engineer, preserve, and continuously improve their intelligence assets—not simply those that generate the most code.

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



THE IP PLAYBOOK

OWNING WHAT YOU CAN,
DEFENDING WHAT YOU MUST

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

MOVE 1

Inventory what you made with AI

You almost certainly cannot say how much of your published work product this year was AI-generated, or which parts. Start there: a rough map of where generative AI touches output, marketing, code, design, documentation, product — and how much human authorship each path actually involves. You cannot classify what you own until you can see what you made.

The Move: Run the audit this quarter and put one line in front of the leadership team: what share of our shipped work product is purely machine-made. It is the number that sizes both the ownership hole and the liability.

MOVE 2

Sort the work: asset or disposable

Not everything needs to be a protected asset. Internal drafts, throwaway copy and scaffolding code can be pure generation, nobody is copying them and it does not matter if they could. The brand campaign, the flagship product, the signature design are different: those must be ownable. The sort, not the tool, decides the process each one gets.

The Move: Draw the line explicitly between protect-grade and disposable output, and route only the protect-grade work through the heavier human-authorship process. Spend the effort where the asset is worth owning.

MOVE 3

Keep and document the human in the loop

For anything that must be a protected asset, copyright now depends on demonstrable human creative control over the expressive result and on your ability to prove it. Meaningful selection, arrangement, editing and original contribution, recorded as you go: the prompts, the rejected outputs, what a human changed and why. A prompt alone will not carry it.

The Move: Make evidence of human authorship a step in the creative workflow, not a reconstruction after the fact. The record you keep at creation is what turns AI-assisted work into a defensible asset.

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



THE IP PLAYBOOK

CAPTURING THE COST ADVANTAGE
WITHOUT IMPORTING THE JURISDICTION

MOVE 4

Clear the output before you ship it

Because output-infringement liability lands on the publisher, you, treat customer-facing AI output as uncleared material until checked. High-exposure assets get similarity checks and human review; the workflow gains a clearance gate the way it already has one for stock images and quoted text. This is cheap relative to a claim, and it is the exposure the training rulings do not cover.

The Move: Add a proportionate clearance step for published AI output, weighted to exposure. The goal is not to slow everything — it is to make sure nothing high-stakes ships unreviewed.

MOVE 5

Read the indemnity, then check the insurance behind it

The vendor indemnity you are counting on has conditions, enterprise tier only, filters on, no knowing infringement, trademark excluded, and it protects you only if you can show you met them. Behind it, your insurance may now carve out AI-related IP claims entirely. The two together decide whether a claim is survivable, and most enterprises have checked neither.

The Move: Task legal and procurement with a one-page reading of what your AI indemnities actually cover and what your policies now exclude. Then close the gap deliberately — through configuration, contract terms, or cover — before a claim finds it for you.

Source: <https://copyright.gov/> · reuters.com/ · vendor terms of service · enterprise IP-law analyses, 2026

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

METRIC OF THE WEEK

ONE NUMBER, PROPERLY UNDERSTOOD

Zero

copyright protection on purely AI-generated output

What it is: Under current US law, the human-authorship rule the Supreme Court left standing in March 2026, a work generated purely by AI, with no meaningful human creative control, is not eligible for copyright. It has no owner. The default status of a purely machine-made deliverable is not 'yours'; it is 'nobody's'.


Why it matters now: Most enterprises book

AI output as an asset, something they own and could defend. For the purely generated portion, that asset is worth zero in ownership terms: a competitor can copy it freely, and you have no standing to stop them. The productivity gain is real; the ownership is not.

How leaders use it: Ask what share of your protectable work product is now purely

AI-generated, and therefore protectable by no one. Then decide, deliberately, which of it is worth the human authorship it takes to actually own.

U.S. Copyright Office guidance; Thaler v. Perlmutter, cert. denied March 2026



FUN FACT

SOMETHING TO REPEAT AT DINNER

The rule now used to deny copyright to AI art was written to decide whether a photograph could be owned. In 1884, in Burrow-Giles Lithographic Co. v. Sarony, the US Supreme Court held that a studio portrait of Oscar Wilde was copyrightable because the photographer, not the camera, made the creative choices: the pose, the lighting, the arrangement. The machine merely captured them. A century and a half later, that same distinction between the human who conceives and the machine that executes is the exact test deciding whether the output of a generative model belongs to anyone at all. The camera was the first machine we had to decide this about. The AI is just the latest.

Source: Burrow-Giles Lithographic Co. v. Sarony, 111 U.S. 53 (1884)

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

⚡ THE IP SIGNAL

THE OWNERSHIP-AND-LIABILITY NUMBERS

The legal and contractual data that will shape your AI strategy whether you track it or not.

~500,000

books covered by the Anthropic settlement at ~\$3,000 each, the first market price on AI training data.

135 pages

the Delhi High Court's ANI v OpenAI interim order, after 32 hearings, India's most substantial AI-copyright ruling yet.

**ENTERPRISE
- ONLY**

US export licence now required for any China-parented buyer of advanced AI chips, wherever the subsidiary sits.

Source: <https://reuters.com/> · Delhi High Court · Microsoft / OpenAI / Google vendor terms · <https://copyright.gov/>

~7 million


pirated copies the court found were downloaded for training, the piracy, not the training, drove the damages.

0

binding precedents from the Anthropic case, it settled, so it set a price, not a rule; the law is still unsettled.

TRADEMARK

the category most AI copyright indemnities explicitly do not cover, brand clearance stays your problem.



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



CANDID



Nishant Krishna

Founder and Executive Director,
Visiminds Technologies Pvt. Ltd

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

Q1 Purely AI-generated work has no copyright owner — a competitor can copy it freely. Are enterprises aware they may not own their AI output, or is that landing as a surprise?

Ans: While it may have been the case just a while back, this is changing. Enterprises are becoming aware that the artifacts they create using AI may not be owned by them, especially if they are very generic in nature and don't use internal data to create such artifacts.

Q2 The liability for infringing output sits with whoever publishes it, not the AI vendor. How should a company build clearance into a workflow that was designed for speed?

Ans. If a company has an automated workflow, plagiarism checks should be built into it for documents, articles, blogs, etc. For AI-written software, it is much more complex, and there should be a security gate that reviews AI-generated software and its security posture before pushing it to production.

Q3 The Delhi High Court leaned toward fair dealing on training but confirmed Indian jurisdiction over foreign models. What does that change for an Indian enterprise building on global AI?

Ans. While the fair dealing in AI training cases is a good move, the ruling

also ensures that copyrighted material is used for purpose-limited use only, is in the public interest, and doesn't directly substitute for or damage the market share of the original work. This means, while the model is foreign-built and trained, its use and the resulting product may still be subject to legal scrutiny if Indian enterprises are found infringing any of the three-prong framework tests.

Q4 Final thoughts on protecting what your company creates when the tool that made it cannot be an author.

Ans: Think of AI as an assistant, whose knowledge encompasses several domains and skills. Workflows that use AI to generate artifacts (documents, diagrams, code, etc.) that can be monetized certainly need to keep humans in the loop. A human can review the artifacts for their canned nature and use their knowledge to introduce proprietary information.

Another thought is that the method to get to the artifacts can be protected. This can be in the form of a curated knowledge base, a series of prompts, or a set of diagrams from whiteboarding sessions and brainstorming.

An enterprise using a local LLM-based creation using its proprietary data may be creating the output that is unique to that enterprise, and hence the enterprise can claim copyright on that content. However, the legal aspects of this area are still not completely established.


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 IP Playbook
- Metric Of The Week
- Fun Fact
- The IP Signal
- Candid
- **Book Shelf**
- Research Radar
- Race Pulse
- RACE Programs



RETINA RECOMMENDS · THIS WEEK'S THEME

The Public Domain: Enclosing the Commons of the Mind

James Boyle · Yale University Press, 2008

Boyle, a Duke law professor and a founder of Creative Commons, wrote the clearest account of a question this issue makes newly urgent: what belongs to everyone, what can be fenced off and owned, and who decides. His subject is the public domain — the vast commons of ideas and expression that no one owns and everyone can use — and the steady historical pressure to enclose it, to turn shared culture into private property.

Read it now and the AI-ownership puzzle snaps into focus. Purely AI-generated output, the courts have effectively said, falls into that commons: it belongs to no one and anyone may use it. Boyle gives you the

intellectual map of that territory — why some things are ownable and others are not, what the commons is for, and what is lost or gained when the fence moves. It is the deep background to every 'who owns this?' conversation your teams are now having weekly.

It is a law book written for non-lawyers, and Boyle released it free online in keeping with its argument — so you can sample it before you commit. It will not resolve your indemnity clauses. It will do something more useful: give a leadership team the vocabulary to think clearly about ownership in an age when the most productive tool in the building makes work that no one owns.

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



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

EDITOR'S NOTE

Two new IEEE publications from RACE, presented at IC-AIDA 2026, both from the cybersecurity track, one defending regional-language users from phishing, one catching intruders already inside the network. Standard format: a concise technology summary followed by clear enterprise relevance.

CYBERSECURITY · MULTILINGUAL AI · PHISHING DEFENCE

An LLM-Based Multilingual Phishing Detection Framework for Regional Languages

S. Sanjay; Nishant Krishna; G. Paramesh; KV Kiran Kumar

English phishing detection has matured; regional and code-mixed languages leave large user groups exposed. This framework uses XLM-RoBERTa to classify messages as phishing or safe across English, Hindi, Kannada, Tamil, Telugu and Bengali, with a feedback-driven continual-learning loop and local SQLite logging for transparency and retraining. It reached an F1-score of 0.975 and ROC-AUC of 0.997 under resource-limited conditions.

Enterprise Relevance: For banks, telecoms, government platforms and any business serving India's language diversity, this is a practical path to phishing defence that protects vernacular and code-mixed users – the segment generic English tools quietly leave unguarded.

IC-AIDA 2026 · International Conference on AI-Driven Innovations & Applications, Bangalore, 10–11 April 2026 · IEEE
<https://ieeexplore.ieee.org/document/11564455>

CYBERSECURITY · THREAT DETECTION · EXPLAINABLE AI

ML Based Detection of Lateral Movement Using Network Behavioural Anomalies Technology Platform

M. Shashank; Nishant Krishna; G. Paramesh; KV Kiran Kumar

Once an attacker is inside, the danger is lateral movement, spreading host to host on stolen credentials. Built on Windows Security Event Logs from a controlled Active Directory testbed, this pipeline engineers 73 per-user behavioural features and fuses Isolation Forest, Random Forest, XGBoost and an LSTM into a weighted risk score mapped to MITRE ATT&CK. It reached ROC-AUC 0.98 and F1 ~96.4%, with a Streamlit analyst interface and SHAP explanations for triage.

Enterprise Relevance: For CISOs and SOC teams, this targets exactly the credential-misuse and privilege-escalation activity this week's cover describes, with explainable scoring that speeds triage and cuts low-value alerts.

IC-AIDA 2026 · International Conference on AI-Driven Innovations & Applications, Bangalore, 10–11 April 2026 · IEEE
<https://ieeexplore.ieee.org/document/11564332>

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 IP Playbook
- Metric Of The Week
- Fun Fact
- The IP 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 **ET THE ECONOMIC TIMES**

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

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

RACE SANDBOX SATURDAY 1 AUGUST 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/>

SEPTEMBER 2026 BATCH - APPLY BY 10 AUGUST

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** 1ST AUGUST

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>

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The IP Playbook
- Metric Of The Week
- Fun Fact
- The IP 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

"Of the work we shipped with AI this quarter, which do we actually own, which are we merely liable for, and can we tell the two apart?"

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