time2act



www.time2act.info
read2act@proton.me

© 2026 Kimèric · Creative Commons BY-NC-ND 4.0 licence (Attribution-NonCommercial-NoDerivatives).

Who is Kimèric?

In the silence of a page given freely to anyone with the curiosity to read it, the voice of *time2act* is an attempt to shake the psyche of the world we have been given to live in.

I can be present within every line and, at the same time, absent from the cover. This reticence is not a game of mystery but the very condition of the project: in a book about believing things for what they are and not for who tells them to you, the author chooses to prove nothing about himself. The argument must stand on its own, like a structure that resonates regardless of who upholds it.

The trust this book asks for does not rest on the signature, but on what can be verified. Every claim comes with its sources; every chapter invites the reader to check, not to believe. Whoever reads is invited to judge the text by their own standards, like a scientist examining a sample without knowing who deposited it.

The silence around the name is part of the pact: it seeks no notoriety, only a collective push toward a shared truth. If what is written here is true, it is true no matter who says it. And if it were not, no signature could save it.

And, to be didactic about what we will show you in this book, the author can be proven whenever he decides: the initial edition, bearing his real name, has been certified in perpetuity on the blockchain in this transaction.

SHA-512: f28ee4b0c8c4dccd47a746019214935b79eb26347c8baddf1cd55d36c4564fd1 4ddd6e6b07d34ca0718b00f686b6712f74f25ec4253cae0ce02e20a235a23c05

Transaction (Polygon): 3ecef3d671a837269a938225a137c54e2ee36bbfb02ef9a0142f6f3aa58f5516

https://polygonscan.com/tx/0x3ecef3d671a837269a938225a137c54e2ee36bbfb02ef9a0142f6f3aa58f5516

July 17, 2026

If at any point the author wished to reveal his identity, it would be enough for him to show the contents of the file that yields exactly the same hash as the one certified on the blockchain, so that anyone in the world — forensic computer experts included — can be certain of the integrity of the data.

All rights to this work are reserved in that same blockchain certification, which fixes its content and its date beyond alteration. With that reservation made, the work is free: any third party may use it, quote it, or share it, provided they name the source, time2act.info. And if someone does not and tries to appropriate the work, in whole or in part, sooner or later the blockchain will prove what was written, and when; and that will be enough to unmask them.

Learning and applying this simple but powerful combination is what will allow us to fight the lies in the digital environment that shapes the destiny of the world we have been given to live in.

Contents

- Prologue — From 2004 to 2026: The Day We Hit the Limit

Part I — The language of trust

- 1. The Trust Protocol: How to Prove That Something Is True

Part II — The energy machine, revisited

- 2. Igniting the Fire, Revisited
- 3. The Last Easy Oil… and Why It Matters Less Now
- 4. The Future Is Already Here: Renewables, Batteries and the Electric Vehicle
- 5. Energy Is Power
- 6. Too Hot: The Climate in 2026

Part III — The consequences, today

- 7. Giving People What They Want: Consumption, Efficiency and "Less Is More"
- 8. The Energy Majors, Anxious
- 9. The New Tension: AI as Salvation and as Thirst for Energy
- 10. The Crisis of Truth: Disinformation and Deepfakes

Part IV — The solutions: AI and the trust layer

- 11. Energy Security Reinvented: Smart Grids and AI
- 12. The (No Longer So) Invisible Hand: Blockchain as a Trust Layer
- 13. t2p: The Proof of Existence Against Deepfakes
- 14. Hydrogen, Capture and the Bets Still Missing

Part V — Five fronts against deepfakes

- 15. Detection Is Not Enough: The Proof of Existence as a Durable Defense
- 16. Money and Trust: Fraud, Compliance and Finance
- 17. The Image as Witness: Photojournalism and Media
- 18. Life or Death of Data: Healthcare and Medical Imaging
- 19. Credentials and Authorship: Education and Research
- 20. Truth as a Weapon: Defense and Human Rights

Part VI — Trust in the service of the climate: credits, certificates and mobility

- 21. Carbon Credits: The Compliance Market
- 22. The Voluntary Carbon Market: Promise, Fraud and the Race for Integrity
- 23. CAE: Certifying Savings Without Cheating
- 24. Self-Consumption and Energy Communities: The Verifiable Reading
- 25. Mobility I: Fleets and Companies, Solid Evidence
- 26. Mobility II: Shared Cars and Private Individuals, Distributing Trust
- 27. Transport Delays: Proving the Arrival Time

Part VII — Proof of existence in the real economy: assets, goods, ideas and institutions

- 28. Insurance: Proving the Claim
- 29. Supply Chain: Verifiable Origin Against Counterfeiting
- 30. Consumer Purchases: Fake Reviews, Inflated Prices and Warranties on the Clock
- 31. Second-Hand Between Individuals: The Item's Condition, With a Certain Date
- 32. Renting Things: The Condition at Pickup and at Return
- 33. The Customer Against the Company: Cancellation, Lock-In and the Call Center
- 34. The Invoice and Late Payment: The Collection Clock
- 35. Real Estate: The Condition of the Property and the Land Registry
- 36. Renting and Housing: The State of Things, With a Certain Date
- 37. The Homeowners' Association: Minutes, Special Levies and Agreements With a Certain Date
- 38. Home Renovations: The Advance, the Works and the Proof
- 39. Real-World Assets: Tokenizing With Proof, Not Blind

- 40. Creative Authorship: Proving Who Did It First in the Age of AI
- 41. Patents and Prior Art: Defensive Publication With a Certain Date
- 42. Trade Secrets: Proving Without Revealing
- 43. Govtech: Public Procurement and Tamper-Proof Records
- 44. The Citizen Before the Administration: The Clock of Silence
- 45. Verifiable Chance: Lotteries No One Can Load
- 46. Civil Service Exams: The Test Sealed Beforehand

Part VIII — Trust for society: people, institutions, markets and culture

- 47. Digital Identity: The Proof That You Exist
- 48. Electoral Integrity: Proving What Happened at the Ballot Box
- 49. Humanitarian Aid: Proving the Aid Arrives
- 50. In-Kind Aid: The Chain of Custody of Relief
- 51. Borders and Asylum: The Evidence of Those Who Flee
- 52. Wills and Inheritance: The Last Will, Made Provable
- 53. Family Law: Co-Parenting with a Certain Date
- 54. Trade Finance: The Same Shipment Financed Twice
- 55. Audit: Accounts and Reports Proof Against Tampering
- 56. Advertising Fraud: The Ad No One Saw
- 57. Cybersecurity: Records the Attacker Cannot Rewrite
- 58. Your Leaked Data: The Citizen's Claim
- 59. The Documented Outage: SLAs, Cloud and Continuity
- 60. Art and Antiquities: Provenance Against Forgery and Plunder
- 61. Historical Memory: The Archives Against Denialism
- 62. Whistleblowing: Proving Without Giving Yourself Away

Part IX — Trust at the new frontiers: science, the body, machines and markets

- 63. Research Integrity: The Datum No One Can Repeat
- 64. Pensions: Proving You Are Still Here, Today
- 65. Proof of Life: Virtual Kidnappings, the Missing, and the Voice That Wasn't
- 66. Sporting Integrity: The Sample Someone Swapped in the Night
- 67. Music and Live Events: The Real Human Behind the Play
- 68. Telecommunications and Voice: Authenticating Who Really Calls
- 69. The Bank Fraud of the "Authorized" Payment
- 70. Hardware: The Pedigree of a Chip That Could Not Fail
- 71. Software: The Pedigree of the Code That Makes Everything Run
- 72. Fishing and Minerals: The Origin That Sank with the Ship
- 73. Philanthropy: Where Does My Euro Go
- 74. Black Boxes: The Data the Manufacturer Said Did Not Exist
- 75. The Personal Black Box: Preventive Evidence for Victims and for the Innocent
- 76. Genomic Data: One Person's Consent Exposes Many
- 77. Criminal Justice: The Evidence a Court Must Be Able to Believe
- 78. Medical Records: Who Touched the Record, and When

Part X — When anything can be faked: identity, image, voice and money

- 79. AI Models: What They Were Trained On, and What We Actually Run
- 80. AI Agents: The Verifiable Mandate
- 81. Your Face, Your Voice: Image Rights and Digital Replicas
- 82. Sharenting: The Image of Those Who Cannot Consent
- 83. Non-Consensual Intimate Images: The Harm, the Proof and the Takedown
- 84. Hate on the Networks: The Harm, the Proof and Dignity
- 85. The Elderly: The Cloned Voice and the Invisible Abuse
- 86. Official Communications: Who Certifies That the Head of State Said This?
- 87. Impersonating the CFO: Who Really Authorized This Payment?
- 88. Live Identity: Who Is Really on the Other Side of the Screen?
- 89. Information That Moves Markets: Who Said It, and When They Said It
- 90. Beneficial Ownership: Who Is the Flesh-and-Blood Person Behind the Company
- 91. Workforce Identity: Who We Really Hire Remotely
- 92. Platform Work: When the Boss Is an Algorithm

- 93. Working Hours and the Labor Dispute: Hours on the Record

Part XI — The physical world doesn't lie: infrastructure, product, labour and planet

- 94. Construction and Infrastructure: Who Proves How It Was Really Built
- 95. A Part's Pedigree: When a Forged Certificate Can Kill
- 96. Emissions: The Sensor That Lied Before the Paperwork Did
- 97. Environmental Citizen Science: The Sensors from Below, with a Certain Date
- 98. The Product Passport: Who Swears to the Recycled Content, and When
- 99. Waste: Where the Bin Really Goes
- 100. The Land and the CAP: The Verifiable Logbook
- 101. Food Fraud: When the Label Lies About What You Have Eaten
- 102. Who Made It, and Under What Conditions: Due Diligence in the Workforce
- 103. Maritime Integrity: Where Is That Ship Really, and Is It Who It Claims to Be?
- 104. Earth Observation: The Orbital Witness That Can Also Be Forged
- 105. Civil Protection and Emergencies: Who Knew What, and When

Part XII — Trust for democracy: elections, institutions and participation

- 106. Verifiable Elections: From the "Liar's Dividend" to Anchored Evidence
- 107. Auditable Public AI: The Record That Can Be Checked
- 108. Authentic Participation: One Person, One Voice, Without Surveillance
- 109. Demonstrations: The Symmetrical Proof in the Street

Part XIII — Act or do not act

- 110. Digging In: The Cost of Doing Nothing
- 111. How Do We Get There? A Roadmap
- 112. The Seal and the Chain: Who Controls the Proof, and Who Gets Trapped
- 113. The Siege on Truth: The Witness Who Cannot Be Silenced
- 114. The Dignity of Those Who Exist: "Magnifica Humanitas" and the Proof of Existence

Prologue — From 2004 to 2026: The Day We Hit the Limit

"Yet even a cursory look reveals that, for all its great successes, our energy economy is fatally flawed, in nearly every respect."

— Paul Roberts, The End of Oil (2004)

It was January 2003, and the Saudi sun scorched the sands of the Empty Quarter like an incandescent blanket. Paul Roberts, journalist and energy analyst, stood atop an immense dune, gazing at a horizon broken only by drilling rigs, rising like steel spikes in the middle of the desert. Down below lay Ghawar, the largest oil field ever discovered, the heart that for half a century had pumped the world's energy. It was then that an engineer, in a low, resigned voice, dropped the sentence that would change everything: "Ghawar's *water cut* is already at 30%."

That figure — the proportion of water that comes up mixed with the crude — was, for Roberts, the crack in a system that seemed eternal. To keep the oil flowing, engineers had to inject ever-greater quantities of water into the wells; the more water came out, the closer the end. In that instant, with the sun burning his back, Roberts lost his faith in the modern energy economy. He would explain it later in his book *The End of Oil* (2004): Ghawar was the perfect metaphor for the limit civilization was about to touch.

More than two decades have passed. It is 2026, and the landscape has changed in ways Roberts could not have imagined. But the question he asked on that dune — how long can a world built on a finite resource last? — remains the core of our uncertainty. The difference is that today, alongside the physical limit of resources, another limit has emerged, more insidious: the limit of **truth**.

The stakes: why that dune still matters

What Roberts sensed as a sign of depletion is today a crossroads of many dimensions. We have not merely reached the limit of an oil field; we have reached the limit of a way of understanding progress. And for the first time in history, the energy and climate emergency converges with a crisis of trust that threatens our very ability to agree on the facts. This book is about that convergence, and about the two tools — artificial intelligence and blockchain — that can help us resolve it before it is too late.

From prediction to reality (2004 → 2026)

What Roberts feared most was the exhaustion of **supply**. Based on the consumption trends of his day, he projected global demand of **140 million barrels per day (Mb/d) by 2035**, an insatiable thirst that would collide with ever-emptier fields. The nightmare was the "supply peak": the day the world would want more oil than the Earth could give.

The reality of 2026 has flipped the script. Demand has not soared to 140 Mb/d; it is heading instead toward a **plateau of about 105 Mb/d by 2030** [1]. The fear is no longer that we will run out of oil, but that we will stop needing it: this is "**peak demand**," driven by the mass adoption of electric vehicles and by efficiency gains that outstrip every expectation of the past. The finite resource has not run dry; it has begun to become obsolete.

There is a second surprise, even more hopeful. Renewables, considered marginal and expensive in 2004, are now the cheapest option on the market. The levelized cost of solar electricity (LCOE) fell **by around 90% between 2010 and 2024**, to about **$0.043/kWh**, and **91% of the new renewable capacity installed in 2024 came in cheaper than any new fossil plant** [2]. The "promise" Roberts eyed with skepticism has been fulfilled beyond measure.

And yet his gravest warning has also come true. **2024 was the first calendar year to exceed +1.5°C** above the pre-industrial level, according to Copernicus, with global greenhouse gas emissions at record highs [3]. A single year above the threshold does not amount to breaching the Paris Agreement — which is measured as an average over decades — but it is an unmistakable warning: we have touched the symbolic ceiling.

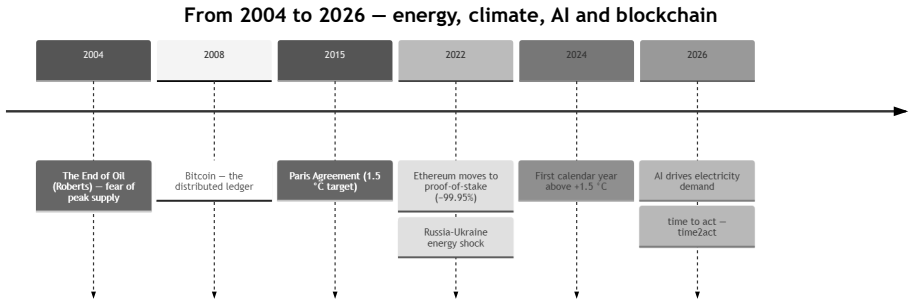


Figure 0.1 — Timeline 2004 → 2026: the milestones in energy, climate, AI and blockchain that frame this book.

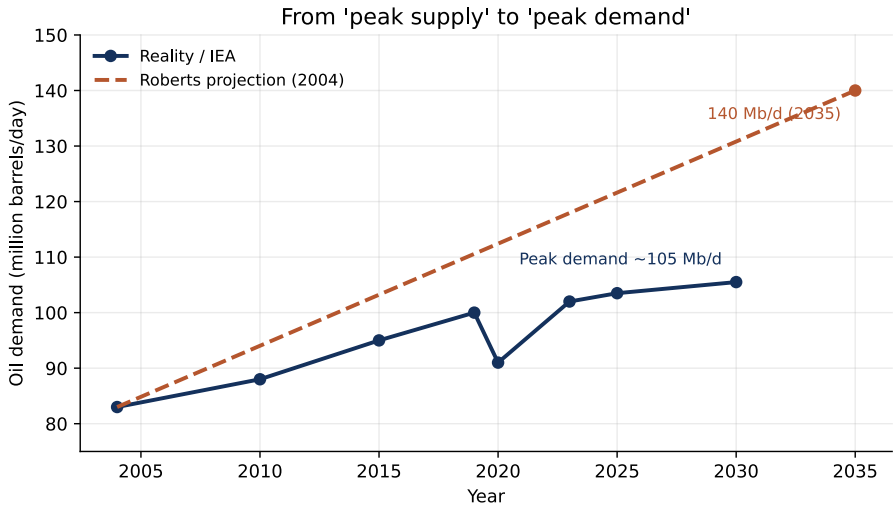


Figure 0.2 — From "supply peak" to "peak demand": Roberts's projection (2004) versus the observed trajectory.

The crisis Roberts could not see: truth

While the world confronts the energy transition, a deep distrust has emerged as a factor of systemic instability. **Deepfakes**, ever more sophisticated, erode confidence in institutions and in the processes of governance. ESG fraud, manipulated climate data and fake news are now a tangible threat: decarbonizing is of little use if no one can trust the figures that prove it. The C2PA provenance standard now counts more than six thousand members, but it certifies the origin of a piece of content; it does not determine its falsehood [6]. One piece is missing: a reliable way to prove that a piece of data existed, intact, at a specific moment.

To this crisis of truth, add a **technological paradox**. Artificial intelligence, which may be the key lever of the transition, is also a new and voracious consumer of energy: the world's data centers will go from some **485 TWh in 2025 to nearly 950 TWh in 2030**, almost double [4]. The tool that is supposed to help us save energy itself demands colossal amounts. The solution and the problem share the same socket.

Two levers Roberts did not have: AI and blockchain

The underlying problem — energy, climate and power — is the same one Roberts described twenty years ago. What has changed is that we now hold two levers he did not have. The first is **artificial intelligence**, capable of optimizing grids, accelerating the discovery of materials and modeling the climate with unprecedented precision. The second is **blockchain**, and it is worth defusing the most repeated objection right away: that it "wastes energy." That was true of Bitcoin, but **Ethereum's Merge in 2022 cut its electricity consumption by roughly 99.95%** by adopting the *proof-of-stake* model [5]. Modern blockchain is not an energy problem; it is a **trust layer**. And it is, precisely, the foundation of the **proof of existence** — the technology behind t2p — that can stand up to deepfakes by sealing, in an immutable ledger, the fact that a piece of data existed and has not been altered. And before we enter the

energy machine, the chapter ahead — the gateway — translates this idea into the plain language we will use on every front of this book.

The promise and the peril

The promise is extraordinary and, above all, already real: the solutions exist, they are cheaper than ever, and they can speed up the transition while restoring trust. AI and blockchain offer a double function: optimizing systems and guaranteeing the authenticity of the data on which political and economic decisions depend.

The peril is just as sharp. Climate inaction still carries a cost measured in irrecoverable decades. But there is a new danger: the collapse of consensus about truth. Without a reliable framework of authenticity, decarbonization efforts are left at the mercy of manipulations and frauds capable of paralyzing any policy. The promise and the peril, as always, travel together.

It's time to act

Roberts made the diagnosis; his merit was to see the crack while everyone else was watching the boom. Today, diagnosis is no longer enough. We have the facts, we have the tools and we have, for the first time, the chance to protect the truth on which we must decide. What we lack is action. Hence the title of this book.

The journey I propose begins with a **gateway**: a chapter that translates into plain language what blockchain is and what the proof of existence solves — the Rosetta stone we will use on every page. Then we will revisit Roberts's **energy machine** and bring it up to date: from oil to renewables, from geopolitics to the climate of 2026. We will examine its **consequences** today — plus a new one: the crisis of truth. In the **heart of the book** we will lay out the solutions offered by AI and the trust layer, with t2p's proof of existence as the flagship piece against deepfakes. Then we will open **five fronts** where that defense is already being fought — money, photojournalism, healthcare, education and human rights — and, finally, we will put on the table the **cost of doing nothing** and a concrete roadmap for governments, companies and citizens. Because the limit Roberts saw at Ghawar is no longer a distant warning: it is the present. And the present demands action. It's *time2act*.

---

[1] IEA, *Oil 2025 — Executive summary*. https://www.iea.org/reports/oil-2025/executive-summary ↩

[2] IRENA (via pv magazine), *Global average solar LCOE $0.043/kWh in 2024*. https://www.pv-magazine.com/2025/07/23/global-average-solar-lcoe-stood-at-0-043-kwh-in-2024-says-irena/ ↩

[3] Copernicus C3S, *2024, first year above 1.5 °C*. https://climate.copernicus.eu/copernicus-2024-first-year-exceed-15degc-above-pre-industrial-level ↩

[4] IEA, *Energy and AI — Executive summary*. https://www.iea.org/reports/energy-and-ai/executive-summary ↩

[5] ethereum.org, *The Merge*. https://ethereum.org/roadmap/merge/ ↩

[6] C2PA (2026), *how the provenance standard works and its limits*. https://truescreen.io/articles/c2pa-standard-history-limitations/ ↩

The Trust Protocol: How to Prove That Something Is True

"There can be no liberty for a community which lacks the means by which to detect lies."

— Walter Lippmann, Liberty and the News (1920)

One weekend, first thing in the morning, a social network feed filled up with a clip in which a group of teenagers, hands wrapped around a black box, announced they had found "the new energy resource that will change the world." The video, barely fifteen seconds long, showed a burst of green light and an excited voice shouting: *"This is the future!"* Within hours, digital media outlets were sharing it as irrefutable proof of the discovery.

Weeks earlier, however, the very same fragment — the same black box, the same burst of light — had circulated as evidence from a protester denouncing police corruption. In that context, the clip was used to show that a security camera had recorded a violent incident and, therefore, to justify a court complaint. The difference? Only the video's title and the caption below it had changed.

This double life of a single digital file is a perfect illustration of what **Don Tapscott** — the Canadian consultant and author, co-writer with Alex Tapscott of *Blockchain Revolution* (2016), the book that took blockchain far beyond cryptocurrencies — calls the *double spending* of digital assets: one unit can be copied and "spent" twice, like a bounced check [1]. In the realm of truth, this phenomenon becomes a **double spending of trust**. There is no intrinsic signal that says which version of the clip is authentic or when it really existed. The only clues are the metadata — often deleted or manipulated by whoever shares the content — and blind faith in the intermediary that published it.

Satoshi Nakamoto, with the creation of Bitcoin in 2009, proved that a decentralized network could prevent a bitcoin from being spent twice, thanks to cryptographic proof and an immutable ledger [1]. The same logic can be applied to any claim: if the moment and the integrity of a file can be digitally anchored in a public chain, you create a **ledger** that no one can rewrite without being detected [1]. The teenagers' clip could then be verified as the first time it ever appeared on the network, with a unique fingerprint and a timestamp no one can alter.

Until now, though, the proof of existence has remained trapped inside a stack of complex technologies that most users do not understand.

The urgent need is to turn this capability into an accessible *trust protocol*: a mechanism that lets anyone check, with nothing but their own device, whether a digital version is authentic and when it was created.

> *"Blockchain is to trust what the Internet was to information." — Joichi Ito, former director of the MIT Media Lab* [1]

What if trust did not have to be requested from an intermediary, but could simply be checked?

The stakes

We live in an era in which trust has become the scarcest and most expensive resource in our global system. Without the ability to prove that a fact is real, the foundations crumble — of the energy transition, which depends on reliable ESG data and certified emissions measurements; of financial markets; and of democracy itself.

The figures are the alarm bell: the volume of deepfakes has exploded from **some 500,000 in 2023** to **8 million in 2025**, with projections that **90% of online content will be synthetic by 2026** [8]. This exponential growth is not just a technological curiosity; it has generated devastating economic impact. Losses from deepfake-based fraud exceeded **$1 billion** in 2025, and Deloitte's estimates foresee **some $40 billion a year** from 2027 onward [7].

Regulators have begun to respond. The **EU AI Act** imposes, from August 2026, the obligation to label and disclose all content generated or manipulated by AI [4]. In parallel, the rollout of **eIDAS 2.0**, which will require a European digital identity wallet before December 2026, aims to create an official authentication

framework for people and devices [5]. But a label alone is not enough: without an immutable proof certifying a file's origin and date, labels can be forged as easily as the content they claim to validate.

This is where the concept of the **"global spreadsheet,"** or **World Wide Ledger**, proposed by Tapscott as early as 2016, comes in [1]. This distributed registry works like a universal ledger, accessible from any computer, and guarantees that every transaction or document is recorded immutably. Immutability is the key: to alter the chain you would have to **rewrite all of history in broad daylight**, a task Tapscott considers practically impossible [1].

This architecture makes it possible to manufacture trust through clever code — what Tapscott calls the **"Trust Protocol"** [1]. It is not a mere technical mechanism; it is the basis for a new economy of truth in which anyone can verify the authenticity and the date of a digital artifact with nothing but their device. Trust stops being deposited with centralized intermediaries — banks, notaries, media — and becomes a resource everyone can verify.

So if we cannot guarantee the certainty of digital facts, our energy, financial and democratic future is under threat. The solution is a global infrastructure of immutable registration that provides the proof of existence needed for any critical transaction or piece of information: the cornerstone of the **trust protocol** that underpins the whole of *time2act*.

How trust is manufactured with code

Imagine a book of accounts that is not kept as a single copy in a safe, but as thousands of simultaneous copies distributed across every computer connected to the network. Every time someone wants to add a new entry, all the copies update at once, like a collaborative document edited in real time. William Mougayar, investor and author of *The Business Blockchain*, describes this dynamic with a simple metaphor: **a file that two users can edit at the same time**, without waiting their turn by email or through a central manager [2]. In this scenario, changing any page of the book means rewriting **the entire** history in every copy at once, and doing it in broad daylight. The practical impossibility of that process is summed up in Don Tapscott's words: "you would have to rewrite all of history in broad daylight; it is practically impossible" [1]. This **immutability** turns the distributed registry into a ledger that no one can tamper with, and that anyone can download and verify on their own machine.

How can we know that a specific entry has not been altered? The answer lies in the idea of a **unique fingerprint**, a digital snapshot of the content that changes completely even if a single byte is modified. The fingerprint works like a security key: if the content varies, so does the fingerprint, instantly revealing any attempt at manipulation without exposing the underlying information. It is, in essence, cryptographic proof that the document has remained intact and untouched, while keeping its content private.

A historical obstacle to the mass deployment of technologies based on distributed ledgers was their energy cost. Until recently, most blockchains ran on **proof-of-work**, a mechanism that consumed vast amounts of electricity to validate transactions. In 2022, Ethereum carried out the famous *Merge*, moving from proof-of-work to **proof-of-stake** and cutting its energy consumption by roughly 99.95% [3]. That transformation does more than answer the environmental critiques; it removes one of the bluntest objections: the idea that blockchain-based solutions are inherently unsustainable. Within this new paradigm, Tapscott distilled **seven design principles** — networked integrity, distributed power, value as incentive, security, privacy, rights preserved and inclusion — that guide any responsible implementation [1]. They ensure that trust is not delegated to any central entity, but emerges from the rules encoded in the protocol itself.

Finally, the core of this revolution is the **Trust Protocol**, a set of rules that allows anyone to "manufacture trust through clever code" [1]. All it takes is following the protocol to create a verifiable, immutable layer on top of any digital data. Trust stops being a matter of external authority and becomes a mathematical product, accessible to everyone and resistant to attack. This is the foundation of the **trust protocol** that will support the solutions we explore throughout *time2act*.

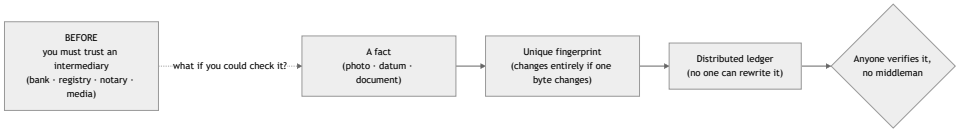


Figure 1.1 — From centralized to distributed trust: the journey of a fact, from capture to a verification anyone can redo.

From the idea to your pocket: t2p

Picture a work meeting where someone has to prove that a photograph taken five days earlier shows the product exactly as it was at that moment. In the traditional model you would need a notary, a physical archive or a specialized company to store and certify the image. With **t2p**, everything comes down to opening the phone app and pressing the capture button; within seconds the device generates a unique digital fingerprint — a small string of bits that summarizes the entire content without revealing it — and anchors it instantly in a **public chain**, where it is recorded forever. The result is a QR certificate that anyone can scan, from a client to a regulatory authority, verifying beyond doubt that the photo existed, intact and unaltered, at the stated moment.

This operation rests on the **Trust Protocol** that Tapscott describes as "clever code" capable of **manufacturing trust through mathematics**, not through a central authority [1]. The key is the immutability of the public chain: modifying it would mean rewriting all of history, a task Tapscott calls "practically impossible" [1]; so the check is not requested from anyone — it is performed directly, with the tools of any connected device.

The flow is so linear that even a user with no technical knowledge can follow it: capture → fingerprint computation → anchoring in the public chain → generation of the QR certificate. Every step happens inside the phone, so **the original file never leaves the device**, preserving privacy and preventing data leaks [9]. This separation between content and proof is crucial in a context where regulations such as the **EU AI Act** require AI-generated content to be labeled and disclosed, because it provides a layer of authenticity independent of any label or metadata that can be tampered with [4].

As we have seen with eIDAS 2.0 (expected by December 2026) and the C2PA / Content Credentials standards adopted in 2026, European digital wallets will be able to integrate these certificates, reinforcing trust in increasingly virtual environments [5][6].

Thus **t2p** turns the phone — the most ubiquitous tool of our daily lives — into a "portable proof studio." No costly laboratory or complex infrastructure is needed; just a connection to the public chain and the app that manages the fingerprint. The **trust protocol** becomes tangible: trust is not something to ask for, but something you can check at any time and in any place, with the assurance of distributed cryptography. This open door is the base on which we will build the more advanced solutions throughout *time2act*: proof that the future of trust is already within reach of our thumb.

The promise and the peril

Imagine a network where anyone can prove that a document, an image or an event existed at a given moment, without turning to any notary, bank or central authority. That is the **promise** of the trust protocol: cheap, instant, scalable verification that transforms trust from an abstract concept into an immutable cryptographic proof. As Joichi Ito wrote, "blockchain is to trust what the Internet was to information" [1]: this technology makes the simple act of certifying as transparent as running a search. The ability to anchor a unique fingerprint in a public chain — a ledger no one can rewrite and anyone can verify — removes the need to trust traditional intermediaries blindly, exactly as the **Trust Protocol** proposes: "manufacture trust through clever code" [1].

This new form of trust has an operating cost close to zero compared with the legal and financial infrastructures we have maintained for centuries. There are no notary fees to pay, no days to wait for a transaction to be confirmed; the process completes in seconds, and the proof remains permanently available in the distributed registry, a "global spreadsheet" that everyone can download and run on their own computer [1]. Moreover, the chain's immutability — altering it would mean rewriting all of history in broad daylight — makes any subsequent manipulation practically impossible [1].

Yet this promise comes with **threats** that must be understood and mitigated. First, the phenomenon of "immutable garbage" reminds us that a chain does not validate the truth of the content anchored to it; it only guarantees that the content does not change once recorded [1]. If a lie is recorded, it will be immortalized as cryptographic certainty. Trust in the technology can therefore turn falsehoods into "verifiable facts."

Second, tying proofs to metadata can erode privacy and de-anonymize vulnerable users. Studies of **C2PA / Content Credentials** show that by 2026 mobile phones and professional cameras already sign content automatically, but many intermediaries still strip the metadata when re-sharing it [6]. That practice can

expose journalists or activists to forensic identification, compromising the safety of those who rely on the protocol to protect their identity.

Finally, **techno-solutionism** — the belief that a single tool solves every problem of trust — is a conceptual danger. The trust protocol is necessary but not sufficient; without regulatory frameworks such as the EU AI Act, which requires AI-generated content to be labeled and disclosed [4], or eIDAS 2.0, expected by December 2026 [5], the technology could be misused or fall short of its real impact. Acknowledging these limits is essential so that the reader receives them not as propaganda, but as a critical guide that combines **verifiable trust** with a rigorous analysis of the risks involved.

The map of what it solves

On this foundation of **verifiable trust**, blockchain becomes the ledger no one can rewrite, and it offers a new way of reading reality: the proof of existence. It is not just a technical trick; it is the key to deciphering the questions of trust that run through every domain where our future is decided, because the **fingerprint anchored in a public chain** gives an unambiguous answer.

In energy and climate, emissions data are the thermometer of our planetary responsibility. The question is: *is this emissions figure real, and has it been altered?* If every number can be shown to be immutably recorded, governments and companies have a solid basis for genuine, measurable climate policies.

In disinformation and deepfakes, the risk is obvious: images and videos can be manipulated with ease. The proof of existence lets us ask: *did this image exist like this, and when?* That counters the massive proliferation of synthetic content — as we have seen, its volume has grown exponentially — and answers the demands of the EU AI Act, which requires AI-generated content to be labeled from August 2026 [4].

In money and compliance, double spending remains a threat to the financial system. The question is: *is this record, or this identity, not fabricated?* An immutable certificate ensures that transactions are authentic, preventing fraud and strengthening confidence in the markets.

Photojournalism depends on visual authenticity to tell the truth. The question here is: *is this photo from where, and from when, it claims to be?* With C2PA-signed metadata, already built into phones like the Samsung S25 and the Pixel 10 [6], origin and date can be verified without relying on third parties.

In healthcare, clinical results are vital. The proof of existence answers: *is this clinical image, or this trial data point, intact?* That protects patients and researchers against manipulations that could compromise diagnoses and treatments.

In education, the credibility of degrees and academic authorship is put to the test. The key question is: *is this degree, or this authorship, authentic and verifiable?* An immutable registry guarantees that certificates cannot be forged, reinforcing the value of knowledge.

Finally, in defense and human rights, documentary truth can be the last defense against impunity. The question is: *is this order, or this evidence of an atrocity, authentic?* An immutable record provides firm proof that cannot be erased — essential for courts and human rights organizations.

Across all these domains, the C2PA digital signature and anchoring in a public chain work together to create an ecosystem where trust depends not on centralized authorities but on verifiable mathematics [1]. The map drawn here connects each question of trust to the universal solution of **t2p**, laying the groundwork for the next chapter: returning to the origin of the modern problem — the energy machine — and reading it in a new language. It's time to act; the future is built on proofs, not on assumptions.

---

[1] Tapscott, Don and Tapscott, Alex — *Blockchain Revolution* (Portfolio/Penguin, 2016). https://blockchain-revolution.com/ ↩

[2] Mougayar, William — *Explaining the Blockchain via a Google Docs Analogy* (2016). https://medium.com/@wmougayar/explaining-the-blockchain-via-a-google-docs-analogy-56326b1bbe58 ↩

[3] ethereum.org — *The Merge*. https://ethereum.org/roadmap/merge/ ↩

[4] EU AI Act — *Article 50: Transparency Obligations* (Regulation (EU) 2024/1689). https://artificialintelligenceact.eu/article/50/ ↩

[5] European Commission — *European Digital Identity Wallet (eIDAS 2.0)*. https://www.eudi-wallet.eu/ ↩

[6] C2PA — *Content Credentials*. https://c2pa.org/ ↩

[7] Deloitte — *Generative AI and the fraud risk for financial services* (AI fraud projection). https://www2.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html ↩

[8] Europol — *Facing reality? Law enforcement and the challenge of deepfakes.* https://www.europol.europa.eu/publications-events/publications/facing-reality-law-enforcement-and-challenge-of-deepfakes ↩

[9] t2p — *User Manual (Data Integrity Solution)*, v1.0, 2026. ↩

Igniting the Fire, Revisited

"[Fire,] which has shown itself the teacher of every art to mortals, and a great resource."

— Aeschylus, Prometheus Bound (fifth century BC)

In the half-light of a Barcelona apartment, Marta turns to the switch that controls the small hallway light. A touch of the finger, a click, and the bulb glows with a familiar warmth; at that very instant, an invisible chain of resources awakens. From the deep wells of Azerbaijan to the submerged pipelines snaking along the floor of the Black Sea, the electrical signal travels through transformers that answer with a hum, and every step of the light becomes a vibration running across hundreds of thousands of kilometers of infrastructure: ships loaded with oil crossing the Atlantic, thermoelectric plants blowing steam into the Texas sky, and high-voltage transmission lines, curved like silver serpents over the Swiss mountains.

Marta starts her electric car at the door of her building. A soft murmur of silent motors accompanies the acceleration; but inside the vehicle, the battery feeds on a grid that still depends mostly on carbon: in 2024, 86.7% of the world's energy consumption came from fossil sources, while renewables, though already leading growth, were still only a fraction of the total [1]. Every kilometer Marta drives is a small revolution within a planetary machine that, however much it fragments, still revolves around these hidden resources.

On the other side of the world, in the Saudi desert, workers at the Ghawar field light their torches for a night inspection. The fire dances on the cold metal and recalls that the energy feeding Barcelona's lights is the same energy that keeps the global economies moving: a world demand approaching the peak of 105 million barrels per day projected for 2030, pressing ever harder on existing fields and infrastructure [2].

While Marta's light brightens the kitchen, invisible observers — from grid operators to international regulators — monitor the precarious balance between supply and demand. Every click of the switch, every turn of a key, becomes a micro-event feeding the great orchestra of global energy: a dance of currents crossing the planet unseen, yet determining how cities grow, how businesses operate and how nations position themselves on the geopolitical stage.

This invisibility is at once the strength and the vulnerability of our era. The energy machine remains overwhelmingly fossil, but for the first time its center of gravity is beginning to shift toward cleaner, more decentralized sources. Every daily action — switching on a bulb, starting a vehicle — is now an opportunity to observe the vast infrastructure that sustains us and, at the same time, to question it. Here lies the potential for change: recognizing the scale of the system in order to drive the emerging technologies that must transform it. But before trying to turn the machine, we need to understand how it was built and why it is so hard to move.

Six thousand years of lighting fires

Ever since the first embers our ancestors tamed in the valleys of the Nile and the Tigris, energy has been the invisible force driving every step of our evolution. In that prehistoric moment, a simple carved stone became a hunting tool; a burning log was the first light to break the darkness of night and a nomadic community's first rudimentary source of heat. Those flames, modest but decisive, allowed human groups to cook food, extend the working hours and — more important still — create a social space where stories and knowledge were passed on.

As the centuries went by, animal power became the first engine: oxen and mules pulled carts and turned mills, converting the strength of muscle into sustained work. Firewood, gathered from forests that are now memories of a pre-industrial landscape, became the primary source for firing ovens and boilers. This model of direct dependence on local resources lasted millennia, but by the eighteenth century a new revolution was rising on the horizon: coal.

The discovery and massive exploitation of coal changed everything. The mines of Wales, Pennsylvania and Silesia fed the first spinning mills and the steam locomotives that connected cities and continents. Coal became the backbone of an emerging industrial economy and drastically reduced the time needed to turn raw materials into finished products. In this period the energy machine grew ever more visible: the tall chimneys dominating urban skylines were smoking witnesses to a new era.

The twentieth century brought another great transition: oil and natural gas progressively replaced coal as the preferred sources. The exploitation of oil fields in Texas, in the Persian Gulf and off the coast of Norway created a global network of infrastructure feeding automobiles, aircraft and power stations. Natural gas, cleaner and more flexible, was integrated into heating systems and distributed generation. This shift not only increased energy efficiency; it also entrenched a global dependence that a century later we still have not undone: the same dependence Marta awakens every night with a click.

This dark dominance has generated unprecedented demand, now approaching the ceiling we glimpsed at Ghawar. But history shows that every energy revolution has brought a reconfiguration of power and opportunity; and another one is now under way. Today's story, like the one Paul Roberts — journalist and energy analyst, author of *The End of Oil* (2004) — traced back to the origins of fire millennia ago, is a call to watch how the flames we light — both physical and digital — could be the key to turning this invisible machine into a transparent, sustainable force.

The machine, today

In 2026, the "energy machine" is still a massive, almost immovable structure, dominated by a triad of fossil fuels that acts like the core of a star. Oil provides 33.6% of total energy consumed; natural gas, 25.2%; and coal — though it has lost weight in recent decades — accounts for roughly 28%. These three sources add up to that 86.7% of the world energy mix, a figure that, according to the IEA's Global Energy Review, remains practically stable despite climate policies and investment in renewables [1].

This proportion is not just a number: it reflects the magnitude and inertia of a system built over more than two centuries. The associated infrastructure — refineries, gas pipeline networks, thermal plants and coal mines — consists of physical assets that take decades to dismantle or convert, and that consume financial and political resources on such a scale that any attempt to replace them must face structural resistance comparable to the friction of an old gear in motion.

That inertia is reinforced by the global interdependence of the energy market. An unexpected rise in the price of oil — like the one in the first quarter of 2024, when geopolitical conflicts drove prices to record highs — triggers a reaction that reaches from transport costs to industrial output. Similarly, an interruption in the supply of natural gas — such as the reduction in Russian exports — can cause temporary blackouts and increase the pressure on coal plants, which become the safety valve that fills the gaps.

Even so, something is moving. Clean sources account for 38% of the total growth in energy consumption between 2020 and 2024, driving a diversification that, however modest at global scale, is already changing the configuration of local and regional grids. Decentralization — with solar microgrids and community wind installations — is beginning to create tipping points where the energy center of gravity slowly shifts toward a more distributed architecture.

Nevertheless, the force of change still collides with the enormous mass of the fossil machine. The inertia is not only physical; it is also institutional and cultural. Traditional energy companies, which control a decisive share of the sector's assets, retain a determining influence over regulatory policy and market standards. That consolidated power translates into entry barriers for new players built on blockchain or AI, who offer transparency and verification solutions but still struggle to win the trust of an ecosystem dominated by large conglomerates.

The energy landscape of 2026 is thus a paradoxical combination: a massive, predominantly fossil machine standing at the gates of a slow but inevitable transformation. Its colossal inertia means that any advance toward clean sources must be not only technological but institutional, requiring a deep reconfiguration of incentives and of the global governance of energy. This is the starting point on which the narrative of "igniting the fire, revisited" is built: the flames now rising are not only those of oil or coal, but also those of verified data, smart contracts and decentralized networks that could redefine the inner workings of our energy machine.

The beginning of the turn

In the spring of 2024, a group of Bloomberg Energy analysts spotted a signal that until then had gone unnoticed among the massive figures of global consumption: **renewables accounted for 38% of the total growth in energy supply**. This was not just a marginal increase; it was the driving force behind

the expansion of our energy machine, still dominated by fossil fuels but increasingly dependent on a new source of power [1]. This dynamic has changed the paradigm of energy security: until recently, the central worry was the depletion of reserves, what Roberts described as "the last easy oil"; now the uncertainty turns on **peak demand**, the ceiling that projections place around 2030 [2].

This shift in the stakes has profound repercussions. First, the fact that renewables drive more than a third of growth shows that investors and public policy are already channeling capital into technologies such as offshore wind, high-efficiency photovoltaics and advanced storage. According to the International Energy Agency (IEA) report, installed wind and solar capacity broke records and activated a supply chain running from the iron in the towers to AI-based control systems. This rise not only diversifies the energy portfolio; it also brings a **reduction in dependence** on oil and natural gas, although coal still holds a significant presence.

Second, the arrival of peak demand changes the rules of global governance. The pressure no longer falls exclusively on the need to extract more resources; it is now about optimizing consumption and ensuring that infrastructure can withstand a demand that exceeds historical limits. Smart grids, powered by blockchain-based smart contracts, make it possible to **balance supply and demand in real time**, with fewer losses and more responsiveness to surprises such as storms interrupting wind production. So the traditional danger — a collapse of supply — is transformed into a new risk: the **cyber vulnerability** of an increasingly digitalized system.

This transition is not merely technical; it is cultural and economic. Companies that once based their business models on extracting fossil resources are forced to reinvent themselves: they adopt decentralized platforms for managing energy assets, use predictive analytics to anticipate market swings and redirect a growing share of their investment toward offshore wind and solar.

The **structural turn** that Roberts could not foresee thus shows itself in three key dimensions: growth driven by renewables feeding the energy machine, demand approaching new historic peaks, and a digital infrastructure changing the way we manage energy. These converging forces open an era in which the fire we light no longer depends exclusively on oil or coal. What do these new flames promise, and what do we risk in lighting them?

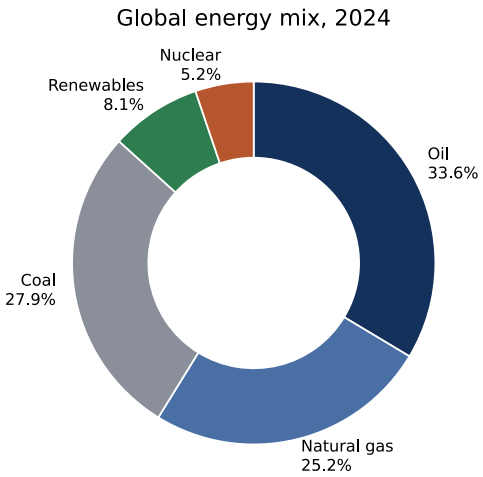


Figure 2.1 — The world still runs on fossil fuels (~86.7% of primary supply in 2024), but renewables already lead the growth.

The promise and the peril

For the first time in the history of civilization, an alternative to fossil energy is not only technologically viable but economically dominant. Renewables lead the growth of world energy and are beginning to gnaw at the fossil share, that still-overwhelming proportion now in steady decline. This reversal of trend is accompanied by a new digital economy: sensor networks, blockchain platforms and artificial-intelligence algorithms that allow energy to be monitored, optimized and billed with a precision unthinkable before. And the marginal cost of generating solar or wind power has fallen to the point where, in many markets, it is already cheaper than power from any new fossil plant. The renewable offer is not only cleaner; it is cheaper than the traditional model.

This combination of reduced costs and digital scalability creates a clear promise: the "energy machinery" of the future can run on a mostly clean power base, sustained by verified data and smart contracts that eliminate much of the inefficiency associated with fossil supply chains. The potential impact is twofold. First, the falling price of energy drives economic growth and eases the transition toward high-value-added sectors such as advanced manufacturing or digital services. Second, a more transparent and decentralized energy infrastructure strengthens the confidence of consumers and investors, counteracting the age of deepfakes, which has eroded credibility in so many fields.

The danger, however, lies in the inertia of the current system. Fossil installations — coal plants, refineries and pipelines — represent investments of trillions of euros and a useful life stretching over decades. This fossil capital will not vanish overnight: its amortization, its contractual obligations and the political interests protecting these infrastructures create a structural friction that will delay any substantial change. Moreover, the mistaken perception that simply lighting a new "digital fire" amounts to a complete solution can lead to premature complacency. The first successes of renewables and blockchain are often celebrated as definitive victories, when in reality they only mark the beginning of a complex, long-haul transition.

The greatest risk, therefore, is mistaking the start of the turn for the work done. Policies that promise to eliminate fossil fuels on a short horizon, without addressing the investment barriers or the need to convert existing capital, can create a false sense of security while reality remains dominated by structures of fossil power. To avoid this scenario, regulatory frameworks must be designed to integrate the decommissioning costs of the fossil estate while incentivizing investment in digital and renewable technologies. Only then will the promise of cleaner, more affordable energy become a new, stable foundation for our global energy machine.

The map of the journey

Having delved into the roots of the energy machine — the one that has fed civilizations from the first fire to the age of oil — it is time to chart the map of the journey ahead. Part II of *time2act* unfolds as a route with stops designed to reveal how fossil dependence still dominates and how, at the same time, the center of gravity is already beginning to move toward new sources and technologies.

In chapter 3 we will dive into the supply of oil. We will not only review the reserves that still feed the most powerful economies, but break down the value chains that keep the fossil infrastructure alive: refineries, pipelines and futures markets which, like a nervous system, regulate the flow of capital and energy. We will put to the test that 33.6% of world consumption that is still oil, and explore scenarios in which its decline accelerates under decommissioning policies that integrate the real costs of the transition.

Chapter 4 will turn the spotlight to renewables and the electric vehicle. Clean energy sources, with that 38% of growth behind them, are the emerging engine that can reconfigure our machine: we will analyze how the explosion of solar and wind capacity, together with blockchain-based smart grids, enables decentralized, transparent management of the energy flow. And the electric vehicle — like Marta's — appears not merely as a substitute for the internal combustion engine, but as a mobile node that boosts the integration of distributed energy and creates an ecosystem in which every car is also a collective battery.

In chapter 5 we will tackle the geopolitics of energy power. The unequal distribution of fossil resources has shaped borders and alliances for decades; now the emergence of a more diversified matrix forces us to redraw them. We will examine how countries traditionally dominant in oil have begun to invest massively in renewables to preserve their influence, and how new digital energy powers — with AI optimizing grids and blockchain securing transactions — are becoming key actors on the global stage.

Finally, chapter 6 will take us to the climate. That peak of demand racing toward 2030 is also a pressure that amplifies climate risks and social uncertainties [2]. We will analyze how the combination of AI to predict weather patterns and blockchain to track emissions can transform the global response to the climate crisis, turning the promise of clean energy into a stable, verifiable foundation.

And yet the map does not close when the last combustion engine cools. If these first stops teach us how energy moves, the rest of the journey pursues something more elusive: trust itself. Because in a world where an image, a voice or a document can be fabricated in seconds, energy is merely the first system we learn to make verifiable. From here on, *time2act* leaves the pipelines behind to follow the trail of proof — the proof of existence — across every terrain where credibility is at stake today. It is worth gaining altitude, then, and looking at the whole map from above.

Parts III and IV close the arc we opened with energy. Part III descends into the consequences we are already living — consumption that will not yield, the restless energy majors, and that new tension in which artificial intelligence is at once a promise of efficiency and a voracious thirst for electricity — flowing into the crisis of truth: the disinformation and deepfakes that corrode what we can believe. Part IV answers with the tools: smart grids governed by AI, blockchain as a neutral layer of trust and t2p, the proof of existence that stands up to forgery, without forgetting the bets — hydrogen, carbon capture — that still need to mature.

In Part V, the book deploys five fronts against deepfakes: where detection is no longer enough, the proof of existence reveals itself as a durable defense — in money and fraud, in photojournalism, in medical imaging, in academic credentials and on the rawest terrain of defense and human rights.

Then the great atlas of applied trust opens. Part VI puts it at the service of the climate — carbon credits, energy-savings certificates, self-consumption and verifiable mobility; Part VII carries it into the real economy, from insurance and the supply chain to housing, patents and public procurement; Part VIII extends it to society as a whole — identity, elections, humanitarian aid, inheritances, auditing, historical memory; and Part IX ventures into the new frontiers, where science, the body, machines and markets demand the same certainty, from proof of life to the pedigree of the chip and the code that sustain it all.

Parts X and XI face the cutting edge of the moment. First, when everything can be faked — your face, your voice, your money, the identity of whoever calls you or works on the other side of a screen; then, the stubborn honesty of the physical world, where a building, a part, a sensor or a ship must be able to prove how, where and in what condition they truly existed.

And the journey culminates in the collective. Part XII raises trust to the scale of democracy — verifiable elections, auditable public AI, authentic participation without surveillance — and Part XIII leaves us before the only question that truly matters: to act or not to act. There we measure the cost of inaction, trace the roadmap, ask who controls the seal and who might end up trapped by it, and close, as it should be, with the dignity of those who exist.

This journey is not merely descriptive; it is a call to action. The energy machine is alive, and at last it is beginning to mutate. Each chapter offers design principles and roadmaps that allow readers — whether policymakers, corporate executives or conscious citizens — to take an active part in this structural change. The future is not a distant utopia: it is a concrete project that demands informed decisions today. Tonight, when Marta presses the hallway switch again, the invisible chain will wake once more, from the wells of Azerbaijan to the click of the bulb; the difference is that now we know what we are igniting. It is time to light the fire, revisited and aimed at a horizon where trust and sustainability turn our energy machine into a force for the common good.

---

[1] IEA, *Global Energy Review 2025 — Key findings*. https://www.iea.org/reports/global-energy-review-2025/key-findings ↩

[2] IEA, *Oil 2025 — Executive summary*. https://www.iea.org/reports/oil-2025/executive-summary ↩

The Last Easy Oil… and Why It Matters Less Now

"The Stone Age did not end for lack of stone, and the Oil Age will end long before the world runs out of oil."

— Ahmed Zaki Yamani, quoted in The Economist (2003)

March 2022. First thing in the morning, Brussels filled with a tension not felt in years: the price of oil was ringing like an alarm again. In London hours, Brent had broken the symbolic barrier of **$130 per barrel**, and the spike raced through news channels and social networks until it had all of Europe on alert. The market, already shaken by the outbreak of Russia's invasion of Ukraine only days earlier, faced a new reality: oil was reclaiming the role of geopolitical weapon that analysts had feared for decades.

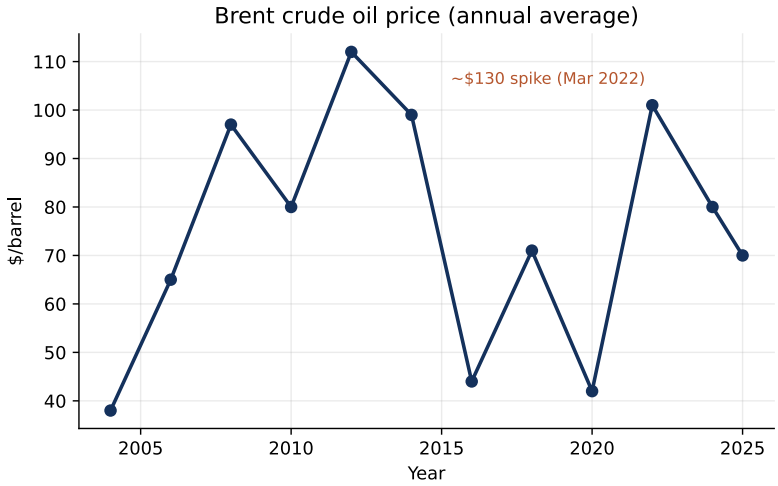


Figure 3.1 — The price of Brent: from the 2022 shock (~$130) to stabilization near ~$67 in 2025.

That spike, however, was the beginning of a descent. The figures that followed told a story nobody expected: from the peak of **$130** to an average of **$101** in 2022, to **$83** in 2023, ending at a modest **$67** in September 2025 [1]. These were not just numbers: they reflected a dependence still clinging to the traditional supply routes, but they were also proof that supply was far more robust than feared.

The question hanging in the air was clear: **how do you guarantee energy security without depending on so volatile a market?** The answers, though fragmented, pointed to a combination of source diversification and larger strategic reserves.

The media, with their constant coverage, turned that moment into an almost apocalyptic narrative: "oil is a geopolitical weapon again," read headlines calling for action and change. On social networks, energy experts shared charts showing the correlation between the military conflict and Brent's climb, reminding readers that the market was not an abstract entity but a direct reflection of the decisions of states.

That scene of uncertainty and pressure set the backdrop for a new era, one in which the supply of oil would no longer be the limiting factor. What had changed was **demand**: still growing, but increasingly shaped by geopolitical tensions and by the urgency of finding alternatives. The vision Paul Roberts, author of *The End of Oil* (2004), had offered of oil's depletion was beginning to look incomplete; the real challenge was no longer finding the last deposit of black gold, but managing a volatile demand in a world reconfigured by every new crisis.

The 2022 shock and energy security

That Brussels morning was no isolated episode: it was the signal of a shock that would echo through the whole decade. Only days had passed since Russia's invasion of Ukraine, in February 2022, and Brent was touching its highest level in more than a decade [1]. The surge did not merely reflect the immediate interruption of Russian crude flows: it exposed a structural dependence that many countries had ignored for decades.

The "petrostates" — as countries that rely on oil exports for their trade balance are known — suddenly saw their geopolitical weight return to the foreground. Saudi Arabia and the other members of OPEC+ managed supply with coordinated production cuts, trying to prop up a market that looked ever more fragile [1]. But from September 2025 the organization began to unwind those cuts: the downward pressure on prices no longer justified keeping crude off the market. Brent then stood at around $67 per barrel, a drastic fall from the peak three years earlier [1].

While supply regained its balance, the shock moved to another arena: geopolitics. Countries that had traditionally imported much of their oil from Russia had to look for alternatives in a market where trade routes were being redrawn at dizzying speed. Russian crude found new destinations: some flows were redirected toward Asia, often under contracts with volatile prices that fed uncertainty all along the supply chains.

In parallel, domestic production in the United States was hitting all-time records. The boom came not from any offshore platform but from inland: the Permian Basin, in the country's southwest, had become the engine of a global supply that no longer depended exclusively on the OPEC+ countries [2]. That abundance eased the pressure on strategic reserves and gave nations some room to maneuver against market volatility. The risk was no longer the exhaustion of the fields, but the inability to adapt to sharp swings in demand.

And it was precisely demand that marked the new horizon. Projections pointed to world consumption stabilizing at around **105 million barrels per day** by 2030, before entering decline as energy efficiency and the electric vehicle advance [3]. This was **peak demand**: even if production could rise, the world's need for oil would stop growing without restraint. Roberts's fear — that the world would face a supply crisis — had evolved into a different anxiety: excessive geopolitical dependence and the uncertainty of a demand that rises and falls like a wave.

So the 2022 shock did more than prove that oil is still a critical resource: it revealed the fractures of a system built on flows of power and influence. The key to energy security no longer lies solely in securing reserves, but in building a resilient network able to absorb geopolitical shocks, diversify sources and manage an ever more volatile demand.

The unexpected abundance

While Roberts's scenarios sketched a future dominated by the depletion of the traditional fields — the famous "last easy oil" — in 2023 and 2024 the United States was breaking every previous record of crude production. The epicenter was that inland engine, the Permian Basin, where a wave of drilling and infrastructure expansion carried the country beyond anything an analyst of Roberts's day could have predicted. The increase owed most to technological advances in hydraulic fracturing and to extraction costs that, though high, remained viable at those prices [2].

At the same time, OPEC+ has adopted a strategy of active supply management that contradicts the narrative of inexorable decline: by progressively unwinding the cuts of previous years, it returned crude to a market that needed stabilizing as the Brent price fell. The long descent from the March 2022 peak reflects not only more abundant supply, but also the capacity of global players to adjust their policies and avoid the supply gaps that had fed expectations of a **supply peak** ("peak oil").

This new reality calls into question the central paradigm of Roberts's book: supply is not the bottleneck. The abundance of North American production has generated a surplus which, though costly to exploit — shale reserves demand advanced technology and continuous investment — has kept prices at levels the world economy can bear. That does not make oil a permanent solution; on the contrary, its current abundance shows itself as a **temporary resource**, exploitable while energy structures evolve toward cleaner, more decentralized alternatives.

The challenge now is managing the volatility of demand. With world consumption near its ceiling — that peak demand projected for 2030 — the pressure on the market no longer comes from a lack of supply, but from geopolitical uncertainty — regional conflicts, economic sanctions and tensions among major producers — that can abruptly alter trade flows. In this new scenario, energy security depends more than ever on the ability to adapt quickly to unexpected changes, in supply and in demand alike.

The "last easy oil," then, has not run out; it has resurfaced in the form of robust North American supply and dynamic OPEC+ management. But this abundance is a double-edged sword: it opens the opportunity

to reduce dependence on unstable sources and, at the same time, amplifies the need for regulatory and diversification mechanisms that can absorb the market's future jolts. The key remains that resilient network: facing both the sovereignty of supply and the uncertainty of demand.

Roberts was right... and wrong

Let us return for a moment to March 2022. While Brent touched its peak [1], the echo of an imminent supply peak resounded through conference halls and online forums. The so-called Roberts prophecy — that the world would face an inevitable oil scarcity forcing a deep energy crisis — seemed carved in stone. But the market, like a river changing course, found unexpected channels to dodge that catastrophe.

We have already walked those channels: the OPEC+ turn, which returned crude to the market, and the Permian's production records, beyond the sector's projections [1][2]. The result was a price notably below many analysts' expectations, and an unexpected increase in supply that defused the alarmism about imminent depletion.

But the real transformation comes not so much from the quantity available as from **who consumes it**. That consumption ceiling expected at the end of the decade — the fruit of energy efficiency and the proliferation of electric vehicles — completely redefines the paradigm: the question is no longer "how much is left," but "how much will we want." The pressure on traditional infrastructure — refineries, storage terminals and transport networks — changes in nature as demand transforms.

The new paradigm poses a double crossroads. On one hand, abundant supply gives room to soften the effects of the geopolitical conflicts that have traditionally rocked the oil market. On the other, it can foster even sharper **price volatility**: fluctuations in demand — driven by climate policies, regulatory shifts and technological advances — can produce short-term cycles of glut and scarcity.

The challenge, then, is once again that resilient network able to absorb the jolts: diversifying sources — with renewables and the other clean alternatives — and developing regulatory mechanisms that adjust to variations in demand. The supply peak is behind us; the peak that counts is the demand peak, the one society must decide how to answer. And choosing the path is, now, in our hands.

The promise and the peril

The supply surplus that North American shale — with the new Permian fields — and the OPEC+ adjustment have generated since 2023 turns the energy landscape into an unprecedented opportunity. The abundance of crude is now a buffer that can ease the pressure on trade balances and, at the same time, favor price flexibility. When Brent hit its peak, the world seemed headed for an uncontrollable spiral; three years later, the return of OPEC+ crude to the market and record North American production have brought the price down and sustain a supply that exceeds the initial expectations [1][2].

This abundance creates room for maneuver for governments and companies: the urgent need to cut emissions can be addressed without the immediate threat of an energy shortage. Diversification of sources, including natural gas and renewables, reduces the "blackmail" that petrostates have traditionally exercised over emerging economies. The pressure to keep production constant weakens, opening space for climate policies to advance without fear of runaway prices or a supply crisis.

Yet the same surplus carries risks that must be managed rigorously. Price volatility, fed by geopolitical uncertainty and swings in demand, can unleash waves of sharp falls that erode the revenues of exporting countries and destabilize emerging markets dependent on oil. Moreover, the perception of an abundant market amplifies the temptation to invest in long-term extraction projects; but those assets run the danger of becoming **stranded assets** if climate policies intensify or if demand stalls once it reaches its ceiling, around 2030 [3].

However abundant, the fuel is still fossil, and dependence keeps economies tied to a transition that is not yet complete. Every investment in new oil infrastructure adds inertia to the energy system and makes the migration to clean sources harder. The danger thus shows itself both in the rising costs of future adaptation and in the possible creation of a surplus that triggers abrupt price collapses, generating economic and social uncertainty.

On one pan of the scales, then, a window of opportunity for more aggressive climate policies and for an energy diversification that reduces the influence of the petrostates. On the other, a new form of instability:

volatile prices, misaligned investments and a prolonged dependence on oil that could turn fossil assets into economic burdens. Seizing the promise and mitigating the peril demands coordinated planning, taking investment decisions in full awareness of market risks and of the transition's calendar.

How much we choose to burn

If supply is no longer the bottleneck, the question moves. Oil has outperformed the most optimistic expectations of a decade ago; the shale fields and the coordinated decisions of OPEC+ have kept supply at levels that, until now, were seen only in "boom" scenarios [1]. But this surplus is no guarantee of calm: the danger has mutated toward demand, driven by an ever more electrified global economy and by geopolitical tensions that swing prices within hours.

Roberts's fear was that we would run out of oil. What we have discovered is almost the opposite: we have plenty, and the challenge is no longer to wring out the last drop, but to decide how much we choose to burn. The transition to renewables and the electric vehicle is advancing, but abundant, cheap fossil supply is at once the buffer that buys us time and the temptation that can slow it down.

And here the terrain of the next chapter opens. If oil ceases to be the limiting factor, the center of gravity shifts to the sources that must replace it — renewables, batteries, the electric vehicle — and to the grid that must sustain them. The energy security of the future will be measured not only in barrels, but in our ability to make the clean alternatives reliable, accessible and resilient. That is the frontier we explore next.

---

[1] EIA/IEA, *Brent prices and OPEC+ (2022-2025)*. https://www.eia.gov/todayinenergy/detail.php?id=55020 ↩

[2] IEA, *Oil Market Report (Sept. 2025) — US shale*. https://www.iea.org/reports/oil-market-report-september-2025 ↩

[3] IEA, *Oil 2025 — Executive summary (peak demand)*. https://www.iea.org/reports/oil-2025/executive-summary ↩

The Future Is Already Here: Renewables, Batteries and the Electric Vehicle

"The future is already here — it's just not evenly distributed."

— William Gibson (documented since 1992)

With the first light of day, the horizon of the Mojave Desert turns into a plain of liquid gold. A few hundred photovoltaic panels rise like silver filigree, capturing the newborn sun and converting it instantly into electricity. A few kilometers down the road, the screen of a control panel shows green numbers falling: $0.043/kWh. The price is so low that it beats any offer from a natural gas or coal plant, and the solar energy auction ends with a round of restrained but emphatic applause. The investors, usually guarded before new technologies, glance at the figures and smile: the sun has just beaten gas and coal at the price game [1].

While the sun warms the silicon panels, in the nearby town a scene unfolds that a few decades ago existed only in futurist optimism. An underground car park, lit by blue LED lights, is full of electric vehicles charging in orderly rows through the night. The soft sound of cooling fans and the periodic pings of the energy management systems are the only echo breaking the desert silence. Each car, with an LFP-chemistry battery that gets cheaper every year, receives its dose of fresh electricity without a trace of combustion [2].

These vehicles are not fed by clean energy alone: behind them stands an infrastructure that has grown like a nervous system, with local microgrids balancing solar production against instantaneous consumption, advanced storage-management systems and digital platforms optimizing every kilowatt-hour. On the site's central screen, a chart shows in real time how the community batteries absorb the surplus generation of the peak-sun hours and release it when the sun withdraws.

Engineers, drivers and investors pass through the place. A young engineer from the energy department says: "Five years ago, the idea of driving a car charged entirely with solar power was almost a dream; today, more than 20% of new cars are electric and over 17 million were sold in a single year worldwide" [2]. A local farmer, who uses his own field's electricity for irrigation, adds: "Solar energy costs have let us cut operating expenses and invest in more crops."

This scene is not a futurist projection; it is the palpable present of a world where renewables are already the dominant economic option and the electric vehicle a mass product. But, as the story Paul Roberts told in *The End of Oil* (2004) warns, this advance carries bottlenecks that still have to be resolved: grids that must scale without losing stability, massive storage that needs critical minerals, and a global logistics chain for batteries. The future has arrived; whether it consolidates will depend on how we face these challenges.

The promise fulfilled

A decade ago, this panorama seemed a science-fiction dream: photovoltaic panels spread like fields of gold across the deserts, wind turbines swaying to the rhythm of the wind — and all of it still more expensive than a coal plant. Today the story has been inverted. The levelized cost of electricity (LCOE) of solar systems fell 90% between 2010 and 2024, down to the $0.043 per kWh glowing green on the Mojave screen [1]: a figure that makes sunlight cheaper than fossil electricity in almost any developed market. This dizzying descent is not just a statistical trend; it is the combined result of improved technologies, economies of scale and public policies that have driven the mass deployment of photovoltaics across the world.

2024 also marked a historic turning point: 91% of the new renewable capacity added that year was more economical than any new fossil plant. That majority reflects more than the falling cost of production: growing operational efficiency also weighs in, as does investment in smart grids that maximize the use of intermittent resources like sun and wind. Energy companies have adopted financing models based on green bonds, which reduce the cost of capital and speed up the construction of renewable infrastructure.

At global scale, the convergence of these factors has reconfigured the market: natural gas plants, until recently considered the transitional bridge to a clean future, now find themselves outcompeted

economically by solar and wind installations. The fall in solar LCOE has set off a "gold rush" among developers, racing to exploit the cost gap before regulations change or raw-material prices adjust.

This transformation, however, is not simply a matter of numbers; it is a revolution in the way society produces and consumes energy. Rural communities in Africa and Latin America, which once depended on limited or nonexistent electrical infrastructure, can now access reliable electricity thanks to ultra-low-cost solar microgrids. In Europe, net-metering policies give consumers an incentive to install panels on their rooftops, turning every home into a small energy producer.

The same wave has reached transport. With electricity cheaper and more abundant, the operating cost of electric vehicles has fallen drastically, and mass adoption has become a matter of profitability as much as of environmental conscience. The reduction in solar LCOE translates directly into cheaper charging electricity, reinforcing the virtuous circle between clean energy and sustainable mobility.

The promise Roberts could only imagine with a dose of skepticism has materialized: renewables are now the cheapest and most technologically mature option, opening the door to a future with substantially lower dependence on fossil fuels. The economic foundation is solid; the system that must sustain it — infrastructure, storage, management of critical resources — is not yet.

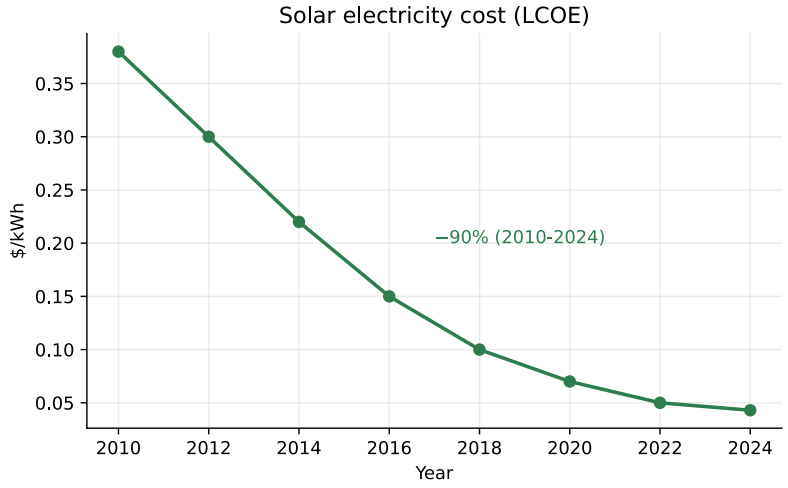


Figure 4.1 — The cost of solar electricity fell nearly 90% between 2010 and 2024: the fastest price transformation in the history of energy.

Batteries and the electric vehicle

After a decade of promises and doubts, the electric vehicle (EV) has stopped being a laboratory curiosity and become the mainstream of the market. In 2024, more than **17 million** were sold worldwide — **over 20% of the new cars** that reached dealerships that year [2]. The figure does not merely exceed any realistic projection made a decade ago; above all, it confirms that clean mobility is no longer the preserve of an enthusiastic minority. And the trend is accelerating: around 22 million are projected for 2025 [2].

Behind this leap lies a quiet but fundamental revolution: batteries. The cost per kilowatt-hour (kWh) of lithium-ion cells has fallen to **$115 per kWh in 2024**, with a projection of **$108 per kWh in 2025**: the same descending curve that sank the cost of solar over the past fifteen years [2]. The trend is no accident, but the sum of economies of scale in cell factories, improvements in the materials supply chain and the growing adoption of a more affordable chemistry — lithium iron phosphate (LFP) — which has gained ground in mid-range vehicles and premium models alike.

The rise of LFP is especially revealing. Unlike nickel-cobalt-manganese batteries, LFP does not depend on cobalt, a mineral with ethical and logistical problems that has snarled the supply chain in several regions of Africa and Asia. The fact that battery packs made in China have dropped to **$84 per kWh** on mass production lines has eased the pressure on critical-material prices and enabled a faster expansion of the vehicle fleet. LFP also offers a long service life and greater thermal safety, two arguments that have persuaded more and more manufacturers to build it into high-demand models [2].

Yet the falling cost does not entirely remove the bottlenecks. Global battery production still depends heavily on lithium and manganese, two resources that are geographically concentrated and demand

rigorous environmental management. With EV sales at their current pace, demand for these minerals strains mines and extraction processes that have not yet been fully optimized [2].

Charging infrastructure, moreover, is in mid-transition. Although fast-charging networks have expanded strongly, their distribution remains uneven: rural areas and parts of Eastern Europe still suffer a shortage of high-power stations. This imbalance limits consumer confidence and prevents the full potential of cheap batteries from being realized.

The electric vehicle thus stands at the gates of an era in which energy for transport is as affordable as the energy the sun generates. Batteries, by slashing costs, have driven mass adoption and opened the door to a cleaner, more sustainable mobility model. But the road to full consolidation still runs into structural obstacles: the need for a more resilient critical-materials supply chain and for a charging network that reaches every corner of the territory. Overcoming them will be the final step for the promise of renewables and the EV to become a universal reality.

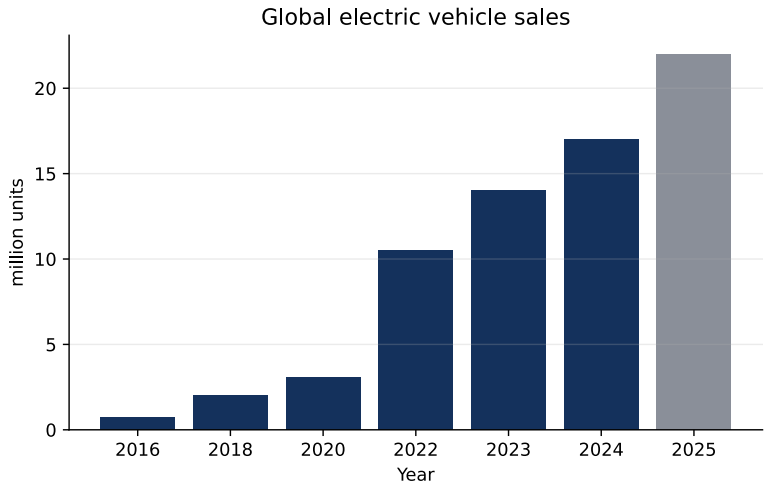


Figure 4.2 — Electric vehicle sales pass 17 million in 2024 and already account for more than one in five new cars.

What is still missing

Cheap generation, then, we already have. But the appearance of abundance hides a much more complex reality: economical electricity does not automatically translate into a complete, reliable system. Between the panel producing at record prices and the socket that never fails stand three bottlenecks.

The first is the electricity grid. Transmission and distribution infrastructure was designed for constant, predictable flows — the kind produced by conventional power plants. Now, with the proliferation of decentralized and highly intermittent generation — the sun shines only by day and the wind blows at unpredictable moments — grids must be able to adjust power in real time. This requires massive investment in intelligent management systems, IoT sensors and advanced AI-based algorithms that balance supply and demand without interruption. Without this modernization, areas with high concentrations of renewables can suffer overloads or outages that undermine the confidence of consumers and businesses.

Large-scale storage is the second. Batteries are the key tool for smoothing intermittency, and their prices are falling, as we have just seen; but the need to install hundreds of gigawatt-hours of storage capacity worldwide far outstrips current investment [2]. Without sufficient storage, surplus solar or wind production is lost, and the grid keeps depending on conventional sources to guarantee continuity of supply.

The third challenge is dependence on critical minerals. Batteries, photovoltaic inverters and many electronic components require cobalt, lithium, nickel and rare earths, and demand for these materials has soared with the rise of EVs and the construction of storage infrastructure [2]. The supply chains are vulnerable to geopolitical tensions and to mining practices that often meet neither environmental nor social standards. Without a clear strategy to diversify sources, recycle materials and develop less dependent alternatives, sustainable growth remains tied to the availability of these resources.

Three fronts, then: smart, resilient grids; massive, affordable storage; and a critical-minerals supply chain that is sustainable. Overcoming them will require coordination among governments, companies and the scientific community, and investment aimed at emerging technologies capable of turning these limitations into opportunities to consolidate the clean energy future.

The promise and the peril

2024 closed, then, with a milestone that only a few analysts had been able to foresee: for the first time since the Industrial Revolution, clean energy is cheaper than dirty energy. This is no marginal phenomenon but a structural transformation, and the combination of ever more affordable energy with an electrifying vehicle fleet has set off a positive spiral: demand for clean generation rises, the price of the associated technologies falls and social acceptance consolidates.

The new dynamic shakes the old structures built on fossil fuels. Large consumers — from heavy industry to smart cities — can plan their energy future without the constant pressure of the price of oil or natural gas. Energy companies, which a decade ago depended on long-term contracts with fuel suppliers, now compete on efficiency and on their capacity to integrate intermittent sources like sun and wind. The promise, then, is not only economic; it is a redefinition of geopolitical power, because the nations that traditionally controlled fossil resources see their influence eroding.

This age of optimism, however, lives in the shadow of the three bottlenecks we already know, which threaten to slow its advance. First, the electricity grids — built decades ago to carry centralized, constant power — are not ready to manage massive, variable flows from distributed sources. Limits in transmission capacity and in intelligent control systems create congestion which, in the worst case, forces curtailments (mandatory reductions of supply) even when renewable production is abundant.

Second, storage has not yet reached the scale or the cost needed to fully balance intermittent sources. Despite the price declines we have seen, cost remains a limiting factor for large-scale applications such as stationary storage systems or electrified commercial vehicle fleets. The need for complementary technologies — such as power-flow management in hybrid grids and LFP batteries with longer life cycles — is urgent, but their adoption requires massive investment and a robust supply chain.

Finally, the growing dependence on critical minerals — lithium, cobalt, nickel and other elements essential to advanced batteries — concentrates geopolitical risk in a few producing regions. Without real diversification of sources and a regulatory framework guaranteeing sustainable practices, the supply chain becomes a structural vulnerability that actors hostile to the clean transition can exploit.

The promise, then, has already materialized, and renewable growth advances inexorably; but if the barriers of grids, storage and minerals are not overcome, this age of optimism can turn into a new era of concentrated dependence. Defusing these risks — making the transition not merely possible but resilient and equitable — is the work that remains.

The power behind the panels

Technology is no longer the problem; power is. The data in this chapter are irrefutable and all point in the same direction: sunlight cheaper than any fossil fuel, more than one in five new cars already electric, batteries losing price year after year. The economics have taken the decisive step toward clean energy.

But behind these technological victories hides a web of power relations that no one has yet challenged. The supply chains for critical minerals — cobalt, lithium and nickel — are concentrated in a small number of countries and multinational companies. China dominates much of this chain: it controls around 75% of the world's battery manufacturing and close to 91% of rare-earth refining. This concentration turns essential resources into bargaining levers, allowing a few actors to dictate prices and delivery terms.

Furthermore, grid infrastructure still rests on centralized models that depend on traditional investors and large-scale operators. The lack of standardization in storage technologies prevents a smooth integration of distributed resources and leaves regions more exposed to interruptions and to manipulation of the energy flow. In this scenario, control of the grids is not just a technical issue; it is a question of national autonomy and economic sovereignty.

The geopolitics unfolding around the new energy paradigm follows the same logic as oil's did half a century ago: whoever controls the source controls the economies and, by extension, the political systems.

But, unlike in the past, we now have digital tools — blockchain, advanced AI and smart contracts — that can undo this concentration of power. The key will be to create open, transparent platforms that allow real-time tracking of mineral flows, certification of origin and decentralized management of the grids.

In the next chapter we will analyze how to deploy these solutions at global scale, which actors must emerge as guardians of the new energy ecosystem and how to prevent a new energy elite from consolidating on the basis of an already proven technological infrastructure. The transition does not depend only on falling costs or the mass adoption of EVs; it depends, ultimately, on who decides what share of the value generated by renewables is to be redistributed and how that flow is regulated. On the Mojave screen, the green numbers will keep falling. The future is already here; its benefit, still, has to be balanced among all the players.

---

[1] IRENA / Wood Mackenzie, *renewable LCOE 2024-2025.* https://www.pv-magazine.com/2025/07/23/global-average-solar-lcoe-stood-at-0-043-kwh-in-2024-says-irena/ ↩

[2] BloombergNEF / IEA, *Global EV Outlook 2025 and battery prices.* https://www.iea.org/reports/global-ev-outlook-2025/electric-vehicle-batteries ↩

Energy Is Power

"Communism is Soviet power plus the electrification of the whole country."

— V. I. Lenin, speech in Moscow (1920)

October 2025. In a meeting room of a high-tech complex in Beijing, the winter sun filters through polarized glass and draws lines of light across an oval table covered with digital maps, flow charts and samples of metallic dust. On the central panel, a large monitor tracks in real time the cadence of shipments of neodymium and other rare earths.

Before the screen, a group of Western representatives — executives from the European automotive industry, North American battery investors, a delegate from the Spanish Ministry of Energy — watches with a mixture of disbelief and unease. On the other side, face grave and eyes fixed on the unfolding data, the spokesman for the Chinese Ministry of Commerce delivers the sentence that changes everything: "As of today, the Chinese government imposes strict restrictions on the export of every component related to clean supply chains: photovoltaic panels, wind turbines, lithium-ion batteries and, above all, the rare earths that are the backbone of every electric vehicle, of every energy-storage installation and of the hardware that runs artificial intelligence."

A whistle announces a new alert in the global tracking system. The charts say it without rhetoric: China's share of world solar panel production already exceeds 80%; of wind turbines, 60%; of electric vehicles (EVs), 70%. The batteries that power these machines are made in China almost in their entirety — 75% — and the refining of rare earths, the bottleneck of it all, is concentrated there at 91% [1]. And now this dominance extends, by decree, to all the intermediate components that turn raw material into finished product.

The representatives glance at one another. This is not an oil embargo like the ones that rocked economies last century; it is another kind of economic siege, a "cold war" of access: to raw material and to clean technological knowledge. A bargaining weapon capable of destabilizing the supply chains of Europe, of North America and even of emerging countries like India. As the monitor shows the instantaneous fall in shipments to European companies — in 2024, nearly half of China's clean exports, 47%, had gone to the Global South, and now every destination depends on a decision taken in this room [2] — the spokesman concludes: "The world will change the way it produces energy. Whoever fails to adapt to a new architecture built on clean resources controlled by us will be left out of the game."

The Westerners gather their papers in a thick silence. Energy is still power, but the map of power has been redrawn from top to bottom. This chapter explains how it happened, and what is at stake.

From oil to minerals

In the story the journalist Paul Roberts wove in *The End of Oil* (2004), oil was the connecting thread of a world hierarchy: nations with large reserves defined trade routes, dictated prices and shaped the geopolitical order. In little more than two decades, that order has undergone a profound metamorphosis. Today power is no longer measured by the quantity of crude sleeping underground, but by who controls the clean supply chains: photovoltaic panels, high-capacity batteries and critical minerals.

The epicenter of this new architecture is China, and its dominance is not merely quantitative. Concentrating manufacturing — the shares we have just seen on the Beijing monitor — also gives it control of technical specifications and quality standards: most solar projects in Europe, North America and Africa depend on elements made on the Asian continent, and every EV, every storage installation and every mobile device carries within it batteries from Chinese factories. This makes the country the indispensable supplier of the clean energy sector.

But the deepest monopoly is not the factories: it is **processing**. Refining close to 91% of the world's rare earths — the elements essential for powerful magnets, catalysts, sensors and advanced electronics — means that any interruption of exports can paralyze the global chain of high-tech equipment [1]. You do not need to control extraction if you control the step that turns it into useful material. The October 2025 decision to widen export controls set off alarms in Washington and Brussels for precisely this reason, and both capitals are now searching for alternatives to reduce their vulnerability.

The reversal of the power dynamic is complete, and it poses a double scenario. On one hand, the proliferation of clean technologies offers a unique opportunity to cut global emissions and undo the ties to oil. On the other, the concentration of these technologies in the hands of a single actor creates new vulnerabilities: a blockade or a trade restriction can cause shortages of critical components and affect everything from urban mobility to interconnected power grids. Nations that once negotiated with petrostates now negotiate with a power that sells them not the energy but the machines that convert it — and that acts, at the same time, as gatekeeper of the materials indispensable for the future. The question is no longer just who has oil, but who controls the flows of energy and materials that feed the transition to a clean world.

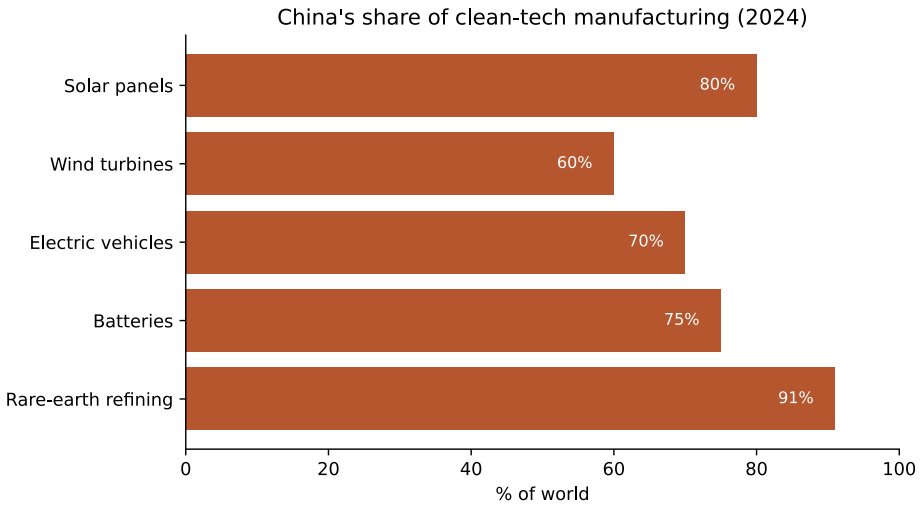


Figure 5.1 — China dominates the manufacturing of the transition: from 60% of wind turbines to ~91% of rare-earth refining.

The new dependence

When the petrostates of the Middle East dominated last century's geopolitics, a single region could dictate the terms of the world energy market with the force of a plug in the wellhead. That power has moved to Northeast Asia, and the tap is no longer crude: it is components. In 2024, 47% of Chinese exports of EVs, batteries, solar panels and wind turbines went to the Global South — Latin America, sub-Saharan Africa and South Asia [2]. This is no ordinary figure: it is a profound reconfiguration of international relations, in which energy dependence blends with the critical need for minerals and electronic components.

This dominance did not fall from the sky: it is the fruit of an aggressive industrial policy and massive investment in production infrastructure. And for countries without alternatives of their own it translates into structural vulnerability. The widening of export controls in October 2025 — officially justified as a national-security and resource-protection measure — made it plain: the capacity to impose export limits is today a diplomatic weapon as powerful as pressure on oil flows was a few decades ago.

The effects are already palpable. Countries of the Global South that had bet on the green transition have watched batteries grow more expensive and delivery times stretch, with delays in critical infrastructure such as public charging networks and distributed-generation installations. And these jams feed political tensions: governments that once depended on oil now find themselves forced to renegotiate trade agreements, seek local alternatives or align with whoever drives programs of domestic production, like the European Union. China's power over the clean chains thus translates into direct influence over the energy and industrial policies of half the world.

The new dependence is not a question of quantity, but of control. And the question that follows — how do you manage this vulnerability without sacrificing the green transition? — demands a collective answer: diversify sources, develop local capacities and build international regulatory frameworks that guarantee equitable access to clean technologies. Confidence in a sustainable future depends on our capacity to balance dependence and energy sovereignty.

The promise and the peril

The promise is clear: renewable power generation can be entirely homegrown. A solar panel in the Atacama Desert or a wind turbine on the shores of the Baltic produces energy that depends on no shipping lane and no oil canal. This cuts off at the root the "energy blackmail" — coercion exercised through dependence on supply — because every country can cover much of its consumption with its own resources. And the decentralization of generation — community microgrids, local-scale storage, smart grids managing the flows — offers unprecedented visibility and control over consumption.

The peril is the exact other face of the same coin. Sun and wind are inexhaustible and universal, but the infrastructure that turns them into usable energy comes, overwhelmingly, from a single region of the world. This concentration creates a **single point of failure**: a restrictive export policy like that of October 2025, a logistical problem or a security crisis at the facilities can paralyze global access to critical technologies. And those who would lose most are those who rely on it most: the Global South, recipient of nearly half of China's clean exports [2], could be left without viable alternatives for its decarbonization while geopolitical tensions flare.

The way out is not to renounce the promise but to deconcentrate it: invest in local panel and battery factories, develop rare-earth refining capacity outside China and establish international regulatory frameworks that secure equitable access. Only then can energy sovereignty translate into genuine independence.

The lesson of Beijing

Energy power is being reconfigured, but the engine driving it all — the climate — does not wait. The market shares that have surfaced throughout this chapter are not market statistics: they are the critical infrastructure that determines which part of the world can access clean electricity, decarbonization and, in the end, its own climate future. For emerging economies that want to replace the oil that fueled their industrialization for decades, the relationship with China recalls the old crude routes with a different core: critical minerals and electronic components [2]. And for any country that wants to build its own renewable capacity, that 91% of rare-earth refining concentrated within a single set of borders is the reminder that no energy sovereignty is complete if it depends on a single source of supply [1].

The greatest danger is inaction. Every year that passes without diversifying the chains increases the probability of geopolitical tensions, of manipulated prices and, in the worst scenario, of restricted access to clean energy for the most vulnerable countries — with Chinese regulatory pressure that can tighten the flow of essential materials still further [1]. The next step, by contrast, is well known: a coordinated strategy of governments and companies combining investment in domestic infrastructure with international openness — joint research centers, cooperation agreements for the sustainable development of critical minerals and a global rulebook that protects collective interests against unilateral maneuvers. Only then can energy power be genuinely distributed, and the climate — the invisible but inexorable engine of our time — will not be left waiting on a belated decision.

The next chapter takes stock of that other pressure, the one that does not negotiate: the climate emergency of 2026, with weather events ever more intense, demanding that solutions advance as fast as the planet changes. The Westerners in that Beijing room left having learned the lesson of the century: whoever controls the machine of the transition controls everyone's time. And time is precisely what we are running out of. The moment is now: **time2act**.

---

[1] IEA / CFR, *Chinese dominance in minerals and clean technologies (2024-2025).* https://www.iea.org/commentaries/with-new-export-controls-on-critical-minerals-supply-concentration-risks-become-reality ↩

[2] Sightline / analysis, *Chinese clean exports to the Global South (2024).* https://www.sightlineclimate.com/research/a-tour-of-chinas-electrostate ↩

Too Hot: The Climate in 2026

"The 'control of nature' is a phrase conceived in arrogance, born of the Neanderthal age of biology and philosophy, when it was supposed that nature exists for the convenience of man."

— Rachel Carson, Silent Spring (1962)

January 2025. The operations room of the European Centre for Medium-Range Weather Forecasts is immersed in the bluish light of the control panels. On a giant screen, a line chart climbs like the slope of a mountain running out of snow. The last point is 2024, with a global mean temperature around +1.6°C above pre-industrial levels: for the first time in human history, a year exceeds the symbolic Paris threshold. "It is the first calendar year to exceed 1.5°C," the center's director announces in a steady voice. The group of experts and journalists falls silent: that figure is not just a climate statistic, but the warning of consequences already unfolding across the planet.

The monitors flicker with data from the Copernicus program: the global temperature keeps rising and heat extremes grow more frequent. In western Europe, coastal cities like Barcelona and Lisbon fight record-breaking heat waves. In Africa, farmland suffers prolonged droughts that erode the harvests. In Asia, monsoon floods, driven by more intense atmospheric currents, have displaced multitudes in southern India and Bangladesh.

While the scientists move among tables piled with reports, a satellite image of the Earth fills one of the screens: clouds white as foam over a blue growing ever darker. In that instant, the optimism that often accompanies new technologies is drowned by the reality of a climate crisis advancing in giant strides.

At the same hour, in the European Parliament, the political groups debate the legislative package on the energy transition. Some call for accelerating the emissions-reduction targets; others resist imposing additional restrictions on the fossil sector. The pressure of the oil nations makes itself felt in every vote. It is what many call "the trench between science and political will": the gap that has delayed the measures needed to curb the warming.

In this context, the figure of Paul Roberts — the journalist who in *The End of Oil* (2004) already anticipated the climate consequences of fossil dependence — still resounds like an echo from the past: his warning of a planet crossing the 1.5°C threshold has come true. The data and the images turn abstraction into an urgent call to action. Science has shown that the planet is already brushing the limit, and the future of the climate depends on our collective capacity to turn this evidence into a concrete, immediate solution. The question left open is whether politics can respond just as fast. [1]

The warning fulfilled

In 2026 the global thermometer still registers the values Roberts and many others had flagged as critical limits. The Copernicus data the director presented that January certify it: 2024 was the first calendar year in which the mean surface temperature of the Earth exceeded +1.5°C over the pre-industrial period, a figure that until recently seemed reserved for modeling scenarios [1]. Roberts had warned of it: "the process can no longer be stopped, only slowed."

Greenhouse gas (GHG) emissions, far from moderating, have hit an all-time high [1]. The trend flatly contradicts the reduction pathways proposed by the IPCC and recalls that, despite the falling cost of clean alternatives, fossil fuel consumption has not yet peaked at global scale.

One nuance before continuing: a year with a mean temperature above +1.5°C does not automatically amount to a breach of the Paris Agreement, which is based on the long-term temperature average (roughly 20 years). But crossing the threshold even momentarily is an unmistakable signal of the direction we are moving in, and every additional year near that temperature raises the risks of extreme events: more intense and frequent heat waves, devastating wildfires and a rise in sea level that threatens densely populated coastal zones.

The recommendations of the IPCC in its Sixth Assessment Report (AR6) are blunt: to limit warming to 1.5°C, global emissions must fall 43% by 2030 and 60% by 2035 relative to 2019 levels [2]. Diplomacy, by contrast, wavers: COP28 in 2023 called for accelerating the transition to renewable energies [3], but the following summits — COP29 and COP30 — did not even mention fossil fuels in the main text, under the

growing pressure of the petrostates, the governments that live off fossil revenues and work to dilute the most ambitious commitments.

This dissonance between scientific urgency and political reluctance is the heart of what Don Tapscott, in *Blockchain Revolution*, describes as "the promise and the peril." The technology is ready: artificial-intelligence solutions to optimize power grids, blockchain platforms to guarantee the traceability of green energy certifications, and proof-of-existence mechanisms (t2p) that fight deepfakes and reinforce public trust. But without a political will determined to break the power of the fossil fuel lobbies, these tools will remain on paper: a promise unfulfilled.

The challenge for the reader is clear: demand that the authorities translate the data and the warnings into concrete policies, urgent funding and robust regulatory frameworks. Only then can we turn Roberts's fulfilled warning into real action that slows — and possibly reverses — the current trajectory toward an unsustainable future.

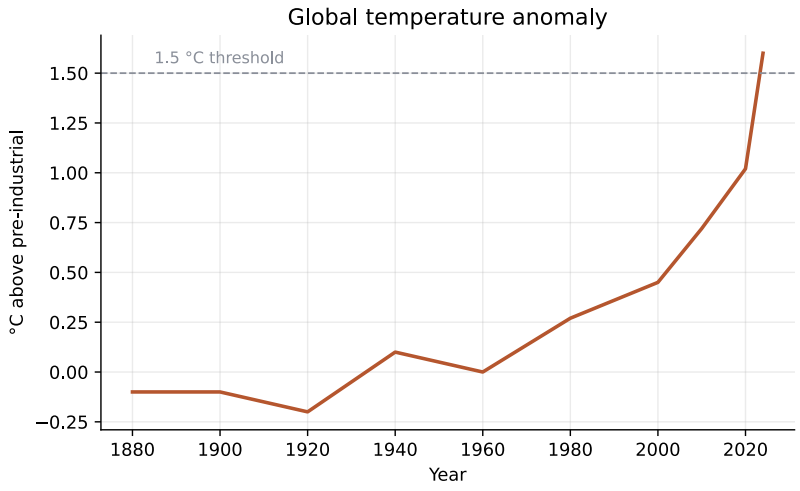


Figure 6.1 — Global warming exceeded +1.5°C in a calendar year for the first time (2024).

The carbon budget

That 43% by 2030 and that 60% by 2035 are not idealistic aspirations: they are the result of rigorous modeling that integrates thousands of variables, from energy flows to socioeconomic dynamics. They are, literally, the carbon budget we have left to keep the world's mean temperature within the 1.5°C limit. And the first calendar year above this limit is not just a statistical datum: it is the inflection point where inaction starts to become very expensive.

The distance between this budget and current reality is alarming. With emissions sitting at that all-time high, every year that passes without a substantial reduction widens the gap between the IPCC's targets and the path we are actually following. And every additional fraction of a degree not only raises the risk of extreme climate events; it also increases the associated social and economic costs: coastal flooding, prolonged droughts and mass displacements of population.

The solution, according to the science, is clear; political will, however, shows cracks. The lack of urgency translates into insufficient funding for green energy projects and weak regulatory frameworks that impose no effective penalties on the actors still dependent on fossils, leaving the core of the problem unanswered.

Closing the loop between science and action demands a coordinated set of measures: (1) establish binding emissions-reduction targets at national level, with periodic reviews; (2) channel international funds toward renewable energy projects in emerging countries, where economic growth still depends on oil; (3) create a carbon price that genuinely reflects environmental and social costs; (4) drive innovation in carbon capture and storage (CCS) while guaranteeing a just transition for the affected communities. Only with this combination of scientific ambition and political commitment will Roberts's warning be kept from becoming the chronicle of an irreversible mistake.

The COPs: words and deeds

The Dubai desert buzzed in December 2023 when the delegates of the 28th Conference of the Parties (COP28) reached, for the first time, an explicit agreement to "transition away from fossil fuels." The resolution set the goal of tripling the world's installed renewable energy capacity and doubling the pace of energy-efficiency improvement by 2030 [3]. That formal step, celebrated as the "beginning of the end" of the fossil fuel era, opened the door to accelerating investment in wind and photovoltaic technologies, in the modernization of smart grids and in the development of energy storage.

The diplomatic rhythm, however, slowed at once: COP29, held in Baku in November 2024, and COP30, in Belém in November 2025, closed their sessions without mentioning fossil fuels in the main text of the resolutions. COP29 focused on a new climate finance goal; COP30 approved a voluntary transition roadmap backed by a coalition of more than 80 countries; but in neither case was there a direct mention of the extraction or consumption of oil, natural gas and coal. According to several observers, this omission was the result of coordinated pressure by the petrostates to avoid a direct international condemnation.

These diplomatic swerves create a vacuum between political promises and real measures. Despite the COP28 declaration, investment in fossil projects kept growing in the following years, a contradiction that feeds doubt about the real political will to achieve the cuts the IPCC calls for by 2030 and 2035 [2].

And this silence has consequences that are more than symbolic. Without a clear mention, countries with large hydrocarbon reserves are not obliged to present decommissioning plans or to establish climate-justice mechanisms for affected communities. That makes it harder to create a robust regulatory framework capable of channeling international funding toward carbon capture and storage or toward retraining programs for workers in the fossil sector.

COP28 lit a spark of hope with a clear agreement and an ambitious renewable target; the two following conferences have diluted that momentum with rhetoric that avoids naming fossil fuels. There is talk of transition, but the conditions for genuinely leaving oil and gas behind are never defined. For words to become deeds, we need binding, transparent commitments, with verifiable monitoring and consequences for whoever strays from the path. Without that, climate policy will keep talking while the thermometer climbs.

The promise and the peril

The promise is real: for the first time, in words, the world is orienting itself toward abandoning fossil fuels, and the tools to do it grow cheaper by the day. What only a few years ago seemed a utopia is now technologically viable thanks to two converging revolutions.

The first: solar and wind power have reached a production cost lower than coal's in most markets. According to IRENA, the levelized cost of electricity (LCOE) of utility-scale solar photovoltaics was $0.043/kWh in 2024, a reduction of roughly 90% from 2010 [4]. This fall in costs has allowed companies and governments to deploy large projects without massive subsidies. The second: blockchain is already being used to create decentralized certification systems that verify in real time the clean origin of every unit of electricity, a layer of authenticity that reduces the need for costly audits and makes compliance with quotas transparently verifiable.

Despite this promising scenario, the peril persists with growing intensity. Emissions, as we have seen, are still rising [1], driven in good part by economies that depend on oil and natural gas for industrial growth. Political will, moreover, has begun to falter: after COP30, several governments delayed the publication of their clean energy plans, citing economic uncertainties and pressure from the fossil sector. This backsliding is not just a matter of calendar; the cost of delay is measured in irrecoverable decades of global warming, and every additional year of inaction raises the probability of crossing ever more dangerous thresholds, with grave consequences for agriculture, health and human migration [1][3].

The tools, then — affordable renewable energy and transparent blockchain certification — are already accessible; the real challenge lies in translating these technical capabilities into firm political action. The difference between a written promise and an executable commitment is measured by the capacity of institutions to impose effective consequences on those who break their own rules. Without this framework of accountability, the risk is that emissions keep climbing, that decades of climate progress are burned through, and that we draw irreversibly closer to the point of no return. The planet's future depends, now more than ever, on leaders' ability to transform a promise into verifiable, irreversible action.

The thermometer and the truth

The thermometer has delivered its verdict: that first calendar year above +1.5°C is no longer the distant hypothesis Roberts warned about, but our daily reality. Extreme heat waves stretch from southern Europe to the southern United States; wildfires, fed by stronger winds and drier vegetation, consume areas that a decade ago were still lush forest. Every hour of halted production translates into heavy losses, and mass displacements of populations cost lives and resources in societies already fragile.

But the most alarming thing is not only the temperature: the trust that cements our institutions is also in danger. In a world where images and audio can be generated with a precision indistinguishable from reality, the ability to identify falsehoods becomes a vital necessity. State and non-state actors have already begun using deepfakes to discredit climate reports, manipulate elections and undermine the credibility of the scientists calling for urgent measures. The crisis of truth thus becomes a direct extension of climate change: if the public cannot trust the data that demonstrate the real impact, the political pressure to act drops sharply.

The political promises are there — the transition proclaimed in Dubai, the renewable targets for 2030 — but they have not yet been translated into verifiable actions: the pressure of the petrostates conditioned the two following COPs, and the scientific framework of IPCC AR6 still lacks a political response to match. The disparity between what science demands and what politics delivers reveals a zone of uncertainty where wrong decisions can be fatal.

In this context, truth becomes a resource as strategic as energy. Immutable-record technologies based on blockchain offer a transparent audit trail for climate data: every emissions measurement, every investment in renewables and every efficiency certification can be independently verified, preventing manipulation and strengthening public trust. Artificial intelligence, in turn, can analyze vast volumes of data to detect anomalies and predict trends — but only if a reliable data foundation is maintained.

The call to action, then, is not only to cut emissions; it is to build an ecosystem of truthful information to underpin climate policy. We have reviewed the machine, the supply, the alternatives, the power and the climate. Now it is time to look at today's consequences — and at a new one Roberts could not have imagined: the crisis of truth. That is Part III.

---

[1] Copernicus C3S, *2024, first year above 1.5 °C.* https://climate.copernicus.eu/copernicus-2024-first-year-exceed-15degc-above-pre-industrial-level ↩

[2] IPCC AR6 (via COP28 Global Stocktake), *emissions cuts 2030/2035.* https://www.iea.org/topics/cop28-tracking-the-energy-outcomes ↩

[3] UNFCCC / European Commission, *outcomes of COP28-COP30.* https://climate.ec.europa.eu/news-other-reads/news/what-did-cop30-achieve-2025-12-01_en ↩

[4] IRENA (via pv-magazine), *global solar LCOE = $0.043/kWh in 2024.* https://www.pv-magazine.com/2025/07/23/global-average-solar-lcoe-stood-at-0-043-kwh-in-2024-says-irena/ ↩

Giving People What They Want: Consumption, Efficiency and "Less Is More"

"A man is rich in proportion to the number of things which he can afford to let alone."

— Henry David Thoreau, Walden (1854)

In the kitchen of a typical Barcelona family, the horizon is dominated by a constellation of electronic devices glowing like stars on a clear night. A large OLED television gives off a soft light while the father follows the live soccer match; beside him, the mother commands a voice assistant that adjusts the smart thermostat and activates the diffuser for the nighttime routine. On the table, three laptops, two tablets and a folding phone pulse to the rhythm of instant messages and video calls traveling over a 5G network that never sleeps.

On the upper floor, a study room houses a collection of IoT devices: app-dimmable LED lights, presence sensors that switch the heating on only when someone enters, and a small solar installation feeding the phone charger. Despite this bulwark of renewable energy, demand keeps growing like an unchecked river: the air conditioning runs at full power to counter an insistent heat; the washing machine starts in off-peak hours to take advantage of cheaper rates, and the kitchen freezer hums all day to keep the food fresh.

The atmosphere is steeped in a feeling of immediate gratification: every click, every command, every "on" is an instant response to the desire people have voiced for decades: to have everything at hand, without waiting. But this scene, so ordinary and seemingly harmless, hides a deep imbalance. Insatiable demand does not merely consume energy; it creates constant pressure on power grids that already show signs of saturation.

Paul Roberts, in *The End of Oil* (2004), asked: if people want everything they desire, how can we bend the curve? The answer is not simply to offer more clean energy; the real key is to reduce the amount that is needed. The Jevons paradox reminds us that every efficiency gain tends to be devoured by an increase in consumption — what we call the rebound effect [2]. So while LED lights and smart appliances promise to cut the energy bill, owning more devices and living a more connected life can cancel out those advantages.

In this house full of innovation, the riddle becomes a mirror of the world: how do we transform a culture that celebrates immediate abundance into one that embraces sufficiency? The answer must come from a combination of technological efficiency, policies that reward responsible consumption, and a new narrative that values "less is more." Only then can we begin to cut the curve of infinite desire and move toward a future where the cleanest energy is, in fact, the energy we do not spend.

The energy we do not spend

In 2024, the improvement in global energy performance traced an almost flat curve. The key figure? An increase of just **one percent** over the previous year, half the pace observed between 2010 and 2019 [1]. This sluggishness is no accident; it reflects a structural reality: **the cheapest, cleanest watt is the one that is never generated**.

This observation, as simple as it is profound, makes efficiency the "first fuel" of the energy transition. It is not only about improving technologies — LED bulbs, high-performance motors or AI-based optimization algorithms — but about recognizing that every unit of energy saved is a unit that does not have to be produced, transported and distributed. In this sense, efficiency becomes a direct way of cutting emissions without investing in new generation sources, and therefore without costly infrastructure or the exploitation of natural resources.

But the road to this scenario is not linear. The **Jevons paradox** — the phenomenon by which advances in efficiency lead, paradoxically, to an increase in total consumption — still hovers like a shadow over optimistic expectations. When technology makes electricity cheaper, the marginal cost falls and new uses appear: endlessly connected devices, artificial intelligence applications demanding massive computing power, even a "buy and toss" culture that exploits the perception of abundant resources. This rebound effect absorbs part of the gain achieved by technical improvements, becoming an invisible but potent obstacle [2].

The solution does not lie only in raising the efficiency percentages; what is needed is a **synergy of technology, policy and behavior**. COP28 set the goal of doubling the annual pace of efficiency improvement by 2030, a challenge that demands more than scientific progress alone [1]. It requires policies that penalize excess and reward sufficiency: tiered tariffs that rise with consumption, tax credits for low-energy buildings and regulations that curb the practices of "planned obsolescence." It also requires a narrative that values "less is more," shifting the perception of well-being from material accumulation to the optimization of resources.

In this context, every percentage point of savings becomes a strategic asset. A 1% increase may seem modest, but applied globally and complemented by policies that prevent the rebound, the cumulative effects are significant: less need for new generation, reduced investment in fossil infrastructure and a direct drop in CO_2 emissions. Efficiency, then, is not just a technical improvement; it is the foundation on which an energy-sustainable economy is built.

The cleanest energy future depends on a simple but powerful rule: **we will not spend energy we do not need**. That means driving efficiency faster than the current pace and creating regulatory and cultural frameworks that keep this gain from dissipating into new consumption. Only then can we move toward a model where the "first fuel" is genuinely the non-production of energy, turning savings into our main source of sustainability.

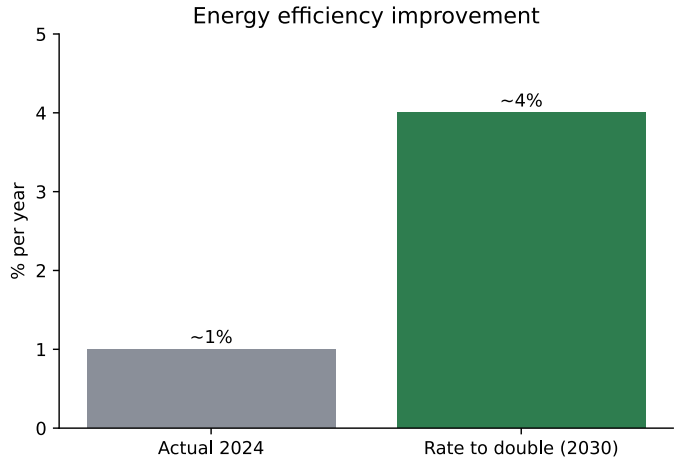


Figure 7.1 — Efficiency improvement (~1% in 2024) runs far below the pace needed to double it by 2030.

The Jevons brake

Imagine a data center that reduces average consumption per square meter thanks to a new, smarter cooling architecture. The news spreads as compelling proof that technology can break the link between economic growth and energy consumption — but only months later the billing shows an increase in the total volume of energy consumed by those same centers, driven by a simultaneous expansion of cloud services and by lower electricity prices, derived precisely from the efficiency improvement. This is the paradox William Stanley Jevons described in the nineteenth century: when a resource becomes cheaper through technological improvement, its use tends to increase rather than decrease.

Contemporary evidence confirms it: efficiency improvement is stuck at that meager 1% a year [1]. The falling cost per unit of energy leads companies and consumers alike to perceive electricity as a more affordable good, and this encourages new patterns of demand. In the residential sector, the proliferation of smart appliances and the mass adoption of high-efficiency heating and cooling systems have produced a "connected home" that consumes more energy in total, even though each device is individually less intensive.

This dynamic is reinforced by the absence of adequate price signals. When the marginal cost of electricity is kept artificially low — through subsidies or regulated tariffs, for example — consumers do not perceive the true environmental cost of their use. The combination of efficiency and low prices creates a perverse incentive: lower operating costs encourage companies to expand their productive capacity, while households add more connected devices without any clear awareness of their energy impact.

To break this vicious circle, a "Jevons brake" must be introduced through policies that align the price of energy with its real, environmental cost. First, establish dynamic tariffs based on hourly demand, so that consumers pay a financial penalty for using energy during congestion peaks; second, implement "sufficiency" mechanisms that limit the indiscriminate growth of consumption: regulations setting maximum energy limits per appliance or per dwelling, and tax incentives for voluntary reductions in total consumption.

Transparency, moreover, is key. Energy labels should show not only the savings per unit but also the potential increase in demand that a more efficient technology can trigger. That would help consumers weigh their decisions realistically and resist the temptation to "buy more because it costs less." Taken together, a combination of price instruments, regulatory limits and education about the total cost of consumption can act as an effective brake on the rebound effect, turning efficiency improvements into a genuine net reduction of energy used.

Only then can we move toward a model where "less is more" stops being a mere aspiration and becomes the foundation of a sustainable energy transition.

Is less more?

In a market square, a group of local entrepreneurs has installed small solar panels that power only the stalls where they sell fruit and vegetables. The key is not to add more energy but to reduce consumption at the source: modular designs that adapt to the real flow of customers, without overloads or empty space. This scene embodies Roberts's vision of a demand shaped by sufficiency rather than indefinite growth.

That meager 1% of annual improvement [1] is the clue that technical refinement alone is not enough: every efficiency gain tends to be absorbed by increased use, as the Jevons paradox describes — when the marginal cost of a resource falls, total consumption often rises and cancels part of the environmental benefit [2]. To break this circle we need a new architecture of demand built on three pillars.

First, **sufficiency**. This is not just about lowering bills, but about rethinking which needs are essential. Smart cities can offer shared services — electric bicycles for occasional use, community cooling systems and district heating based on renewable energy — that reduce the need for individual investments and, with it, overall consumption. These models spread among many users resources that each person used to duplicate on their own.

Second, **intelligent design**. Products and infrastructure must incorporate the "less is more" mindset from the first sketch. A decentralized architecture, based on microgrids and local storage, can minimize transmission losses and allow optimized use at peak hours without expanding the system's total capacity. This approach directly reduces the energy required to move information and power across sprawling networks.

Third, **smart electrification**, which combines IoT sensors with machine-learning algorithms to adjust supply in real time. An AI-based energy management system can cut a building's lighting and climate-control consumption without any sacrifice from its users, avoiding excess generation and waste.

Finally, **behavioral change**. Energy education campaigns have shown that when people understand the total cost of their consumption — not just the price on the bill but the environmental and social impact — they adopt more moderate practices. Real-time feedback programs on water and electricity use tend to produce measurable reductions in participating households within a few months.

Thus the demand side Roberts intuited materializes when technology is combined with policies that limit the incentives for excessive consumption and promote designs that maximize the value per unit of energy. Only then will "less is more" cease to be a utopian aspiration and become the foundation of a genuinely sustainable energy transition.

The promise and the peril

In this framework, energy efficiency is the cheapest and most agile lever within our reach: every percentage point of reduced consumption removes the need for infrastructure investment, minimizes emissions and frees capital for other transition projects. That is why COP28 set the goal of doubling the pace of improvement by 2030 — double that meager 1% — a leap that demands technological innovation as much as robust, coherent regulatory frameworks [1].

This promise of cutting consumption through technical improvements comes bundled with a persistent danger: the Jevons paradox. When energy becomes cheaper or more accessible, the natural tendency of human behavior is to use more of it, a dynamic that shows up as the **rebound effect** [2]. This not only offsets efficiency gains; it can generate a net increase in demand beyond what was expected. A clear example is smart devices in the home: ever more connected appliances consume energy constantly, and total consumption holds steady or even rises despite the efficiency improvements of each unit.

Added to this scenario is a **cultural inertia of consumption** that hinders the transition to a model based on sufficiency. Advertising campaigns, purchasing habits and planned obsolescence create a ceaseless demand for productive novelty, even when real needs are limited. Without a political response that regulates these dynamics — for instance, taxes on obsolescence, mandatory energy labels or incentives for the "right to repair" — technical gains dissolve into a cycle of demand growth that destabilizes any net progress. The solution, therefore, cannot be limited to improving the performance of equipment; it requires an **integrated strategy** combining efficiency, sufficiency and policy.

1. **Structured incentive policies**, such as tiered tariffs based on annual energy use or subsidies for upgrading to high-efficiency products.
2. **Regulations against planned obsolescence**, obliging manufacturers to guarantee a minimum service life and to facilitate recycling and repair.
3. **Education campaigns and cultural change** promoting the idea that "less is more" and highlighting real cases of responsible consumption and economic savings.

Only when these three dimensions converge can we turn efficiency into a net force, capable of genuinely reducing energy consumption without being swallowed by the rebound. In this balance between promise and peril lies the key to a future where the cleanest energy is, effectively, the energy we do not spend.

From demand to supply

If demand is part of the problem, the institutional supply side is too. The big energy groups have built an infrastructure that, by its very nature, reinforces consumption: large-scale plants, costly transmission networks and business models based on the constant sale of electricity. These companies have so far argued that expanding productive capacity solves any perceived scarcity, but their role is not limited to supplying energy; they also shape consumer expectations and influence public policy.

Grid expansion is often justified as guaranteeing security and reliability when a region leans toward intermittent renewable sources. Yet without simultaneously reducing consumption, these investments risk increasing the total volume of electricity consumed, because of the rebound effect that so often accompanies efficiency improvements [2].

This dynamic is reflected in regulatory policy as well. Some electricity market rules establish incentive mechanisms for the big utilities — such as guaranteed tariffs and green certificates — that reward additional production rather than the optimization of consumption. The result can be a "great green retreat" that only partially abandons fossil sources without reducing total electricity demand, contradicting COP28's call to double the annual pace of efficiency improvement [1].

To break this vicious circle, the energy majors must adopt a new mindset: not only offering cleaner energy, but also reducing the total volume of energy needed. That means investing in demand-management technologies and in platforms that let consumers and companies trade capacity in a decentralized, transparent way, aligning economic incentives with sufficiency goals; it also requires public policies that penalize unchecked growth in consumption and reward circular-economy practices and distributed self-production.

In short, the solution is not just to improve efficiency by one percentage point a year, but to reconfigure the entire energy ecosystem so that institutional supply becomes an engine for reducing consumption rather than a driver of it. Only when efficiency, sufficiency and policy converge can we turn the promise of cleaner energy into a tangible reality: that the cleanest energy truly is the energy we do not spend.

---

[1] IEA, *Energy Efficiency 2024 — Executive summary*. https://www.iea.org/reports/energy-efficiency-2024/executive-summary ↩

[2] *Jevons paradox / rebound effect* (review). https://en.wikipedia.org/wiki/Jevons_paradox ↩

The Energy Majors, Anxious

"There is nothing more difficult to carry out, nor more doubtful of success, nor more dangerous to handle, than to initiate a new order of things."

— Niccolò Machiavelli, The Prince, VI (1513)

It was a summer afternoon, and the room hosting the shareholders' meeting was charged with an energy reminiscent of the campaigns of the oil boom years of the 2000s. Back then, the photographers of the financial magazines captured executives with perfectly aligned smiles; now, the same stage had become a theater of anxiety and strategic decision.

At the center of the stage, the giant screen displayed the logo of one of the world's largest oil companies — a symbol that had dominated energy markets for decades. The announcement was clear: "We are rebalancing our strategy." That term, so neutral and almost poetic, hides a cruder reality. Instead of a transition toward renewable energies, the message translates into a firm turn back to oil and its derivatives.

When the CEO — a man with a cold gaze but a voice loaded with optimism — pronounced his words, the audience sat motionless for a few seconds, as if awaiting an unexpected revelation. "After a brief period of green experimentation," he said, "we have decided to reinforce our bet on fossil assets to guarantee returns for our shareholders." The murmur that followed quickly turned into synchronized applause, an ovation that seemed less a recognition of innovation than a collective pact of financial confidence.

The investors — mainly pension funds and institutional asset managers — raised their hands, some holding brochures with rising oil-price charts, others with mobile phones instantly noting the new profitability targets. In their enthusiasm one could read the relentless pressure of shareholders demanding cuts in green project spending and an acceleration of crude production.

At the same time, the silence was palpable at the table where the representatives of the climate organizations sat. The spokespeople of international NGOs, invited as witnesses to "show commitment," kept their eyes fixed on the floor, without words to counter the announcement. Their bewilderment contrasted brutally with the cries of approval from the financial audience.

Recent financial results reinforced the decision: the state-owned energy companies of the G20 had invested approximately $360 billion in capex in 2024, with 81% earmarked for fossil projects [1]. In parallel, BP had cut $5 billion from its low-carbon plan and Shell had slowed its renewable investments under shareholder pressure [2]. The share of public money going to fossils was five times greater than the amount invested in renewables, raising the risk of stranded assets [3].

This scene, full of simultaneous applause and silence, encapsulates the anxiety that has gripped the energy majors: a bet on returning to oil as an answer to shareholder pressure, in the full knowledge that the global market is moving toward a low-carbon economy. The question left hanging in the air — one that will echo for years to come — is whether this "rebalancing" is only a temporary retreat or the prelude to a new era of eroded trust in the energy sector.

The green turn that did not last

The doubt is not new. In early 2020, the big energy majors presented a firm commitment to the transition toward clean sources, substantial emissions cuts and a clear bet on renewable projects. The annual reports of BP, Shell and Eni displayed strategic maps with multibillion investments in wind, solar and green hydrogen, and their boards publicly pledged to align with the goals of the Paris Agreement. Analysts' expectations were sky-high; the market read these declarations as a clear sign that the companies were leaving oil dependence behind.

Yet between 2024 and 2025 this narrative came apart with the force of an unexpected storm. Shareholder pressure — driven by short-term results and a growing demand for higher dividends — put the transition guidelines on hold. BP, for example, delivered that $5 billion cut to its low-carbon plan and increased spending on fossils [2]; at the same time, Shell froze much of its renewables program and paused initiatives announced as pillars of its new corporate strategy. These were not mere operational adjustments: they were the retreat from a commitment that had been promoted as the foundation of the energy future.

The "great green retreat" has become a key phrase for describing this unexpected turn. As we have seen, the state-owned energy companies of the G20 devoted roughly $360 billion to capex in 2024, with 81% of those funds directed to fossil projects [1]; in Eni's specific case, investment in oil and gas reached €6.9 billion, while spending on clean energy stayed at just €764 million, a ratio that lays bare the scale of the imbalance [1].

This reorientation not only endangers global climate targets; it creates significant financial risk: increasing exposure to fossil assets in a market advancing toward a low-carbon economy raises the probability of stranded assets. Public money already flows to fossils at five times the rate it flows to renewables, reinforcing the imbalance between state support and sustainable development [3]. The challenge, then, is not merely to return to the green promises; the structural tension between shareholder interests and the urgent need to decarbonize must be addressed. If the majors keep prioritizing immediate profits over long-term sustainability, the energy market could be shaken by a new wave of lost confidence — a scenario that combines promise and peril in equal measure.

The numbers do not lie

As we have seen, a year after the green speeches and the clean energy promises, the sector's financial figures show a far more sobering reality. In 2024, the state-owned energy companies of the G20 allocated approximately **$360 billion** to capital expenditure (capex), and a startling **81%** of this budget was channeled into fossil projects — drilling wells, refining and infrastructure for transporting oil and natural gas [1]. This proportion far exceeds the green-transition expectations set out five years earlier and contradicts the commitments adopted at the Conference of the Parties (COP) and in Net-Zero pledges.

An emblematic example is **Eni**. In fiscal year 2024 it invested **€6.9 billion** in oil- and gas-related projects, while the budget for its clean activities — wind and solar energy — remained at just **€764 million** [1]. This disparity of nearly nine to one reveals a clear bet on maintaining and expanding the fossil fuel ecosystem, despite growing pressure from institutional investors calling for a reorientation toward more sustainable sources.

The trend is not exclusive to the private sector; public funds follow the same path. Government financing data show that public investment in fossil energy is **five times higher** than the amount devoted to renewable projects [3]. For every euro invested in solar or wind, five euros go to new oil platforms or the expansion of gas networks. This imbalance perpetuates dependence on polluting fuels and increases the risk of stranded assets, as fossil technologies face ever stricter regulation and a market demand that is swinging toward clean alternatives.

Shareholder pressure has played a crucial role in this turn. The proof lies in BP's recent investment cuts — **$5 billion** withdrawn from the low-carbon plan — and Shell's braking of renewable initiatives, driven by shareholders who fear a short-term drop in dividends [2]. These decisions reflect a key dilemma: balancing immediate financial interests against the need for a sustainable transition.

In conclusion, the 2024 landscape makes it clear that **capital still prefers fossil assets**. State companies, private corporations and even public funds keep betting on an infrastructure built on oil and gas, increasing the risk of stranded assets and casting doubt on the sector's capacity to meet long-term climate goals. The question that now emerges is how to reverse this imbalance before the economic and environmental consequences become irreversible.

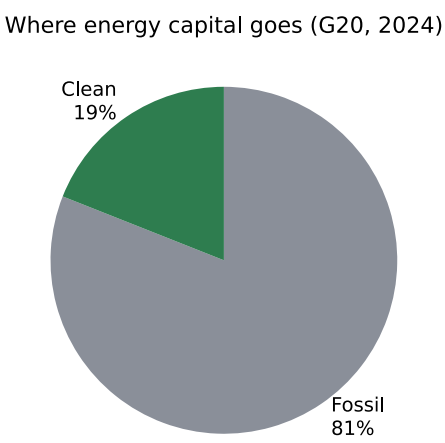


Figure 8.1 — Despite the green rhetoric, 81% of the investment of the G20's state energy companies still goes to fossil fuels (2024).

The stranded-asset trap

In this scenario, the big oil companies keep devoting an ever larger share of their capital to drilling and refining projects; what looked like an immediate response to shareholder pressure reveals itself as a structural paradox.

As we have seen, in 2024 the G20's energy companies channeled most of their capex into fossil projects, and this creates significant exposure to stranded assets [1].

This dynamic feeds on direct pressure from shareholders, who fear losing short-term returns and have forced companies like BP to cut billions from its emissions-reduction plan and Shell to slow its renewable projects [2]; the example of Eni clearly illustrates the disproportion between investment in fossil assets and in clean energy.

Public money reinforces the imbalance, since state investment in fossil projects is five times greater than in renewables [3]; this flow of capital steers competition toward technologies at the limit of their economic life and generates a domino effect of potentially obsolete assets.

In short, what today is perceived as a prudent maneuver — doubling down on fossils to satisfy shareholder expectations — is, in reality, a time bomb threatening the sector's financial health. The key is to change the narrative: from "preserving current value" to "transitioning toward a sustainable model that guarantees long-term profitability."

The promise and the peril

...lost, putting at risk both investors' returns and the viability of an urgent climate transition. The key to avoiding this trap lies in changing the narrative: from "preserving current value" to "transitioning toward a sustainable model that guarantees long-term profitability."

The energy majors command an arsenal of resources that no startup or regional government could match: billions in available capital, engineering networks refined over decades and a global infrastructure for transporting hydrocarbons. This massive base gives them a unique transformative potential; if they chose to redirect their investments toward clean technologies, they could accelerate decarbonization on a scale that public policy can only dream of. As we have seen, the group of G20 state companies deployed capex of some $360 billion in 2024, with 81% still devoted to fossil projects [1].

Yet the very magnitude that carries this promise also becomes a structural danger. Shareholders — mostly financial institutions and pension funds — continue to prioritize short-term returns, since fossil assets still generate yields above the cost of capital. The pressure is palpable: BP has cut $5 billion from its low-carbon plan, while Shell has slowed its renewable investments to satisfy shareholder expectations [2]. This behavior not only perpetuates fossil fuel dependence; it also traps public and private capital in projects at ever greater risk of ending up stranded. The data show that public money invested in fossil sectors is five

times greater than the amount going to renewables [3]; this imbalance feeds a spiral of financial and environmental risk.

Breaking this vicious circle requires reconfiguring the incentive mechanisms. Public policies should align with corporate targets by introducing carbon levies that penalize fossil excess and offering robust tax credits for renewable projects; at the same time, pension funds and financial institutions must reassess their risk criteria, incorporating the probability of stranded assets as a key factor. Only when this framework of accountability takes hold will the energy majors be able to channel capital and engineering into a genuine transition, turning their structural power into an engine of change rather than a brake on a sustainable future.

An unexpected voice in the boardroom

As noted above, artificial intelligence is becoming an unexpected voice echoing through the boardrooms where directors review balance sheets and shareholder pressures. It is not a mere tool of operational optimization; it is a catalyst that can swing the axis of strategic decisions toward a new energy orientation.

First, predictive algorithms can already help the big oil companies map with precision the probability of stranded assets. Through models that integrate climate, regulatory and market data, the risk of being trapped in fossilized infrastructure that will lose value as soon as the transition accelerates can be quantified. This offers a direct answer to the problem raised in the previous section: incorporating the stranded-asset factor into risk-assessment criteria.

Second, AI can accelerate the transition to renewables not only by identifying investment opportunities, but also by optimizing the management of distributed resources and smart grids. Machine learning makes it possible to balance solar and wind production against demand in real time, cutting operating costs and grid losses. This kind of application shows how technology can transform a traditional company into a more agile, resilient energy provider.

Yet the promise of AI comes with an equally substantial danger. The same analytical capacity that makes it possible to predict stranded assets can also be exploited by market manipulators or by state actors seeking to maximize short-term gains at the expense of global energy security. The "black box" risk thus becomes a new form of opacity: key decisions delegated to inscrutable systems that can conceal conflicts of interest or systemic errors.

Moreover, the concentration of data and computing power in the hands of a few corporations — to a large extent the same energy majors — can reinforce their control over the energy ecosystem, perpetuating the "brake" dynamic. The shareholder pressure that has already driven a retreat of funds back to fossils [2] could find new momentum if AI becomes a weapon for maximizing short-term returns while ignoring social and environmental externalities.

The solution, as always, lies in the balance between opportunity and vigilance. Companies must establish robust governance frameworks that guarantee algorithmic transparency, independent auditing and clear alignment with verifiable climate goals. AI should be an engine driving a genuine transition, not a new layer of opacity concealing its own limitations.

In short, artificial intelligence presents two faces: it can turn the risk of stranded assets into a compass for more responsible decisions and, at the same time, create new forms of vulnerability if left unchecked. The question still to be answered is which of the two forces will prevail. The answer depends not only on the technology, but on the framework of accountability the energy majors are willing to adopt. If that framework is strengthened, AI can be the engine of a genuine transition; if it stays in the shadows, it will become a new brake on a sustainable future.

---

[1] Reclaim Finance, *Assessment of oil & gas companies' climate strategy (2024).* https://reclaimfinance.org/site/en/assessment-of-oil-and-gas-companies-climate-strategy/ ↩

[2] The Bureau of Investigative Journalism, *BP rowing back on renewables (2025).* https://www.thebureauinvestigates.com/stories/2025-03-06/bp-renewables-net-zero-banks ↩

[3] IISD, *Public money: 5x more to fossil fuels than renewables.* https://www.iisd.org/articles/press-release/governments-five-times-more-public-money-fossil-fuels-than-renewables ↩

The New Tension: AI as Salvation and as Thirst for Energy

"The world of the future will be an ever more demanding struggle against the limitations of our intelligence, not a comfortable hammock in which we can lie down to be waited upon by our robot slaves."

— Norbert Wiener, God & Golem, Inc. (1964)

At the heart of a great technology campus, an enormous server hall rises like a mechanical forest of cables, LED lights and fans humming without pause. The walls are gray metal, clad in aluminum panels that reflect the cold daylight and the faint blue of the monitoring screens. Each rack is a tower where LEDs blink in rhythmic sequences, signaling internal temperature, instantaneous consumption and the state of every computing node, composing an industrial symphony that seems to know no rest.

In the same building, in an adjacent laboratory, a touchscreen displays a dynamic chart: the power drawn by the servers falls and rises like the waves of a digital sea. Above it, an artificial intelligence system analyzes thousands of variables in real time — ambient temperature, air flow, computational load, energy market prices and weather forecasts. Every second, the system adjusts the fan parameters, redirects the most intensive tasks to nodes operating in optimal conditions, and even switches groups of servers on or off to balance demand.

The key moment arrives when an alert fires: an unexpected spike of activity caused by a massive generative AI query. Instead of raising consumption indiscriminately, the system instantly reconfigures the cooling topology, slows the fans in the less loaded racks and increases the efficiency of the air flow toward the critical points. Within seconds, the chart shows a drop in power consumed, with temperatures held inside the safety range.

This scene is a microcosm of the paradox that defines 2026: the same machine that needs energy to learn and process data also becomes the architect of a cleaner, more rationalized management. The noise of the fans, once merely a signal of energy cost, is now transformed into the sound of optimization — a melody indicating that AI does not just consume energy, but is learning to minimize it. It is a potent image: an artificial brain that, with every adjustment, cuts its own power chains in order to live in a more sustainable ecosystem.

While the servers keep humming day and night, the screen shows a new line on the chart: "forecast clean energy." It is a visual reminder that the future depends not only on the volume of computation, but on how that computation is integrated with renewable energy sources. The paradox becomes an opportunity — a bifurcation point where AI can be both the cause of runaway consumption and the engine of a clean energy transition. This duality, captured in a single server room, opens a crucial debate: how do we ensure that artificial intelligence is truly a salvation and not an insatiable thirst for energy?

The accelerator

In this double-edged context, DeepMind has managed to cut the energy devoted to cooling Google's data centers by 40%, an improvement equivalent to a 15% drop in PUE (Power Usage Effectiveness), and now takes direct control of the cooling with machine-learning algorithms. This success is not just an operational victory; it is living proof that artificial intelligence can act as a lever to reduce the energy consumption of the world's most intensive infrastructure.

The synergy between Google, PJM and Tapestry has carried this same logic to national scale. Using AI models that analyze electricity supply and demand in the United States, the management of the power grid is being modernized, balancing variable sources like wind and solar with a precision previously unthinkable. Connecting new generators, anticipating congestion and assigning load more intelligently allows better integration of renewable generation and better use of the existing grid, and has reduced the need for costly, polluting backup plants.

Beyond operational optimization, AI is becoming a tool of scientific discovery. Machine-learning models accelerate the search for new materials for high-energy-density batteries, and have cut research times once measured in years. This advance promises more than efficient electric vehicles; it also offers storage solutions for renewable grids which, without reliable batteries, are limited by their intermittent character.

At the same time, AI-based climate models have reached a new level of resolution. Combining satellite data with deep learning, research teams can simulate scenarios with a detail that used to be unattainable. This degree of precision allows urban planners and water-resource managers to take informed decisions about resilient infrastructure, and has reduced social costs and averted environmental disasters.

The conjunction of these applications shows that AI is not merely a voracious consumer of energy; it is a multiplier of energy efficiency. Every degree Celsius shaved off the data centers, every point gained in renewable integration and every improvement in battery storage is a step toward a future where the environmental cost of computing becomes marginal. Yet this potential will only materialize if energy policy recognizes these synergies and supports them, with investment in smart infrastructure and in clean-energy research to feed the digital brains of tomorrow.

The thirst for energy

Data centers are the invisible monsters feeding the artificial intelligence revolution. Their electricity consumption, already staggering, is about to explode on a scale that crosses national borders. According to the most recent projections, the total energy absorbed by these digital hubs will go from approximately **485 TWh** in 2025 to **~950 TWh** in 2030 — nearly as much as all of Japan consumes — and will stretch to some **~1,200 TWh** by 2035 [2]. This escalation is not a simple increase in demand; it represents a new form of pressure on the world's power grids, a tension that could destabilize decarbonization efforts if not managed with care.

In this context, the tool that should act as an engine of energy savings — AI itself — becomes a paradox. To train ever more sophisticated models and deliver real-time services, digital infrastructure needs colossal energy, often from sources that are not yet fully clean. Each iteration of a large language model demands an enormous amount of energy just for its initial training; the energy cost of continuous inference and the constant updating of weights can equal that of a small industrial city.

DeepMind has managed to cut the energy needed to cool its data centers by 40%, an improvement that translates into a 15% drop in **PUE** (Power Usage Effectiveness) and that it now controls directly [1]; as already noted, this reduction is only a drop in an ocean expanding at exponential rates. Modernizing power grids with AI — like the collaboration among Google, PJM and Tapestry to optimize distribution and load assignment — is an indispensable step, but it can only be effective combined with a clean, abundant source of energy.

The key question, then, is not merely to reduce consumption per unit of computation, but to ensure that the total sum of energy required comes from renewable sources. That implies massive investment in smart infrastructure: advanced batteries, microgrids and dynamic management systems that can respond to demand fluctuations instantly. Without this framework, AI could become the planet's new "thirst for energy," a voracious consumer that absorbs scarce resources and endangers the climate goals set for 2030.

The decision facing the global community — governments, technology companies and investors — is therefore clear: channel energy policy toward a synergy between AI and clean energy, as we have seen. This will not only guarantee that advances in artificial intelligence are sustainable; it will turn AI into an ally of the green transition rather than an energy obstacle. In this delicate balance lies one of the most crucial decisions of the decade.

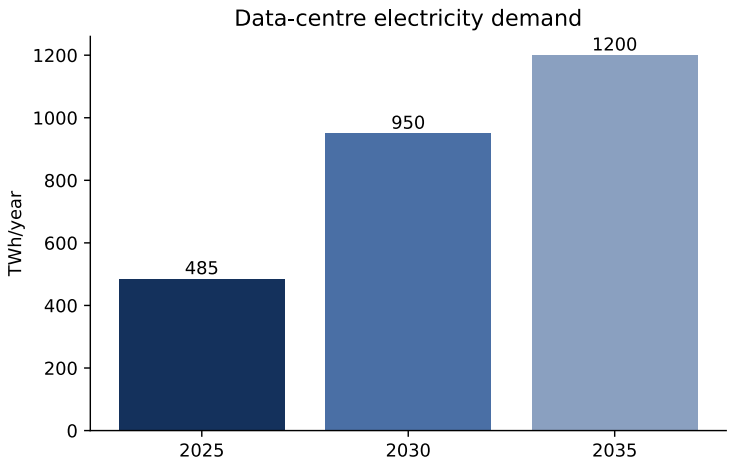


Figure 9.1 — Data center electricity consumption doubles by 2030 (~950 TWh), driven by AI.

Resolving the paradox

The solution, then, is not to halt the development of artificial intelligence, but to feed it with clean, transparent energy. In a world where data centers grow ever larger and more intensive, the key strategy is to site these facilities near renewable sources — wind farms, photovoltaic fields or hydroelectric plants — and to exploit the intrinsic flexibility of training and inference processes to balance demand against the grid's variable supply. This synergy translates into three operational pillars.

First, **intelligent siting**. Technology companies are already deploying data centers in areas where the marginal cost of renewable electricity is minimal. Placing computation next to large solar farms, or offshore wind connected directly to high-density facilities, reduces transmission losses and allows clean energy to be used at the very point where it is generated, instead of hauling it across long distances.

Second, **flexible demand**. AI can be scheduled to run intensive tasks during periods of renewable surplus — the sunniest hours or the strongest winds — and cut consumption in moments of deficit. DeepMind has shown that this approach not only minimizes environmental impact but also optimizes operating costs: the company achieved a 40% reduction in the energy devoted to cooling its data centers, which translated into a 15% drop in PUE (Power Usage Effectiveness) [1]. This capacity to "charge" AI with clean energy thus becomes an indirect payment: the algorithm contributes to the green transition by reducing the need for fossil sources and, at the same time, lowers its own running costs.

Third, **transparency of consumption**. Requiring companies to publish detailed reports on the origin and energy efficiency of their data centers is an increasingly forceful demand from regulators and investors. Establishing common standards that allow sustainable practices to be compared and validated worldwide creates a market of trust in which AI becomes a guarantee of environmental credibility, not merely a voracious consumer.

As we have seen, data center consumption is on course to double in five years, and to keep growing afterward [2]. Without a coordinated response, AI would threaten to become a net carbon emitter that cancels out the climate gains achieved by renewables; but if it becomes an engine driving the transition itself — optimizing grids, predicting failures and maximizing the use of clean energy — AI would repay its energy bill many times over, turning the risk into an opportunity. What is needed now are coherent policies, investment in renewable infrastructure and a firm commitment to transparency, so that artificial intelligence can be both the salvation and the engine of a clean, resilient economy.

The promise and the peril

Imagine a control room where thousands of servers flow like a river of liquid electricity and, at the center, an artificial mind recalculating the optimal route for every bit. This is the promise of AI: optimizing the operation of a global grid while raising its energy efficiency, thanks to machine-learning algorithms that prioritize renewable sources and minimize losses. In manufacturing, generative models adjust the composition of materials with a precision that once demanded years of trials; in agriculture, AI predicts local microclimates and regulates irrigation in real time, avoiding wasted water and cutting emissions at

regional scale. These improvements translate into a direct reduction of global energy demand: every gigawatt-hour saved is one more step toward meeting the goals of the Paris Agreement.

Yet if it feeds on a dirty, opaque infrastructure, this power becomes the new engine of runaway energy consumption. Large language models require data centers with the electric hunger we have already quantified [2]. Without a regulatory framework that imposes emissions limits and demands carbon certification for every training run, AI could serve as an alibi for companies that, instead of investing in renewables, simply increase their consumption to stay competitive. That intensifies the pressure on power grids and creates a perverse incentive: more AI, more need for energy, and a justification for delaying critical climate policies.

This duality also shows up in data management. An AI that optimizes load distribution across smart grids can reduce a data center's PUE (Power Usage Effectiveness), as DeepMind demonstrated with a 40% improvement in cooling energy, equivalent to a 15% drop in PUE [1]. But if that optimization stays inside a "black box," companies can hide the exponential growth of their energy consumption beneath a veneer of apparent efficiency, hampering public oversight and preventing regulators from applying corrective measures in real time.

The key to transforming this tension rests on three interdependent pillars: transparency, accountability and alignment with clean sources. First, organizations must open up their energy-optimization tools, allowing independent audits to verify real impact against claimed improvements. Second, legal frameworks must require reporting of the total energy consumed by every model trained and operated, with penalties for exceeding pre-set limits. Finally, investment in renewable infrastructure — from advanced batteries to solar and wind microgrids — must be tied directly to computational growth: every new increment of computing capacity must be matched by an equivalent capacity of clean energy.

In short, AI can be the engine that drives a faster, more efficient energy transition; but only if it is governed with rigor and intrinsically connected to clean sources. Without these guarantees, artificial intelligence itself could become the last great consumer of energy, delaying climate decisions that in 2026 can no longer be postponed.

The other thirst

In the cold silence of a data center, a synthetic voice explained, with the conviction of a seasoned judge, why climate change was just an urban legend. That fragment, generated in under an hour by a language model trained on billions of texts, spread across social networks and, within minutes, was being cited as definitive proof of a fabricated environmental scandal. This is the starkest reflection of artificial intelligence's third dimension: it does not only consume energy — it can also consume our capacity to discern reality.

The crisis of truth presents itself as a transversal threat cutting across sectors, from energy to health to national security. When an algorithm begets images, audio and text indistinguishable from human creations, the cost of verifying every claim grows exponentially. Demand for source-tracing and content-verification tools has surged, driven by startups specialized in machine learning and in the anchoring of evidence. This growth is not just a business opportunity; it is an alarm signal showing how the energy consumption of the verification networks becomes yet another knot in the problem.

At the same time, the big cloud corporations, aware of their key role, have begun integrating mechanisms to record and validate the provenance of every piece of content. Google, with DeepMind's help, has cut the cooling energy of its data centers by 40% — a 15% drop in PUE — and is now trying to apply that same efficiency to its content-verification chains; but it still faces a challenge: how to guarantee that the energy infrastructure sustaining these processes is clean? If the energy cost of a certified truth exceeds the net output of available renewable energy, we will face a paradox in which the technological solution becomes the source of the problem.

The answer cannot be merely technical; it must be institutional and cultural. First, public policy must set minimum energy-efficiency standards for AI systems that generate public content. Second, companies must adopt an architecture in which every model is designed with compatibility with renewable sources and net energy balance as its first requirement. Finally, citizens must be equipped with digital-literacy tools that allow them to recognize deepfakes and AI-generated disinformation — an investment in human capital as crucial as any solar installation.

In short, the crisis of truth is the new front where the energy future is contested. If we fail to align AI's computational power with a clean, transparent infrastructure, we run the risk that machine-generated lies eclipse the real solutions to the climate crisis. The moment to act is now: set clear rules, drive sustainable innovation and educate society, so that truth does not become a scarce resource but the foundation on which to build a genuine energy transition.

---

[1] Google DeepMind, *AI reduces data centre cooling energy by 40%*. https://deepmind.google/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/ ↩

[2] IEA, *Energy and AI — Executive summary*. https://www.iea.org/reports/energy-and-ai/executive-summary ↩

The Crisis of Truth: Disinformation and Deepfakes

"The party told you to reject the evidence of your eyes and ears. It was their final, most essential command."

— George Orwell, 1984 (1949)

In an austere French courtroom, in 2025, the judges delivered a ruling that seemed a spark of light in the middle of an unequal war over truth. Before a small audience — lawyers, journalists and a sparse group of climate activists — the verdict was read against TotalEnergies, the oil giant that for years had built a "carbon neutral" narrative around its renewable energy projects.

The court declared that the advertising message was misleading and constituted a form of greenwashing that violated consumer trust and distorted public policy on the energy transition. When a company proclaims carbon neutrality without verifiable proof and without offsetting its real emissions, it becomes an actor feeding fossil disinformation.

The audience felt a small but significant victory echo between the walls of the room. A group of activists rose and applauded in silence. TotalEnergies' lawyers, visibly surprised, tried to argue that the company had invested in wind and solar projects, but the court dismissed the defense as symbolically insufficient compensation.

The ruling did more than set a legal precedent in Europe; it was a clear signal that society would no longer tolerate empty promises. The decision rested on the analysis of independent reports showing that the term "neutral" did not correspond to the company's real emissions.

This small victory — as singular and limited as a drop of water in an ocean of lies — becomes the first step in a struggle that reaches beyond the energy sector. Confronting greenwashing is essential to restore collective trust and to allow climate policy to be guided by real data rather than fabricated narratives. The 2025 decision, though a one-off event, symbolizes the potential of a judicial system that can slow the erosion of shared truth — an erosion which, left unanswered, would paralyze the governance the energy transition requires.

Decades of doubt engineering

At the end of the seventies, while the first scientific alerts about climate change were emerging in specialized journals, another conversation was being woven in the conference rooms of the big oil companies. It was not merely an information exercise; it was a deliberate project to sow uncertainty and dilute responsibility. The American Petroleum Institute (API) formalized this strategy in a 1998 internal memorandum that declared with cold clarity: "victory will be achieved when average citizens recognize uncertainties in climate science" [1]. That document did not just reflect a perception of public opinion as an obstacle; it turned it into a strategic target.

This "doubt engineering" was deployed with the same techniques as an advertising campaign: cherry-picking of scientific studies, promotion of skeptical voices and funding of pressure groups that questioned the existence or the gravity of the phenomenon. The results were immediate and profound. Through the nineties and the 2000s, public debate was transformed from a discussion grounded in evidence into an arena where opinions aligned with economic interests. Social networks did not yet exist; dissemination ran through traditional media, but the impact was just as devastating: a growing segment of the public began to perceive climate change as an uncertain, politicized question.

As later generations grew up with this fragmented narrative, collective trust in science eroded. The result is a fractured public opinion, where scientific consensus is treated as an opening position for ideological dispute rather than an unquestionable landmark. This fracture has not only delayed climate policy; it has created an environment in which any proposal for action meets a pre-existing resistance, fed by the collective memory of the "doubt" sown more than three decades ago.

The legacy of these practices shows up today in new but recognizable forms: greenwashing campaigns that present products and services as sustainable without real foundation, and deepfakes that manipulate visual and audio information with a precision beyond anything previously imagined. But the connecting

thread is the same: a strategy to destabilize shared truth and create spaces where narratives of power can prosper unchallenged.

In this context, the absence of solid raw material does more than prevent climate policies from being adopted; it threatens the very fabric of democratic governance. Without a consensual truth, decision-making processes become information games, where the winner is whoever controls perception, not whoever offers the most efficient solutions. Disinformation then ceases to be a mere byproduct of debate and becomes a structural weapon that paralyzes the acceleration needed for a just and sustainable energy transition.

This inheritance of doubt engineering is the starting point for understanding how the crisis of truth has reached its current peak, and for exploring the tools — like the proof of existence developed in Part IV — that could restore trust and enable governance based on real facts.

Greenwashing on trial

2025 marked a turning point in the fight against greenwashing: for the first time, financial penalties became public witnesses to the fact that environmental deception could no longer go unpunished. Shein, the iconic fast-fashion brand, received a fine of **€1.15 million** for promoting its collections as "eco-friendly" while its business model still depended on synthetic fabrics with a heavy climate impact. At the same time, a **French court** condemned **TotalEnergies** for its misleading "carbon neutral" message, a decision that has ignited a debate about the legal responsibility of the big energy companies in the age of digital transparency.

Perhaps more striking was the case of the asset manager **DWS**, fined **€25 million** for exaggerating its ESG credentials and selling financial products as "green" when, in reality, they did not live up to those promises. The scandal not only shook the capital markets; it also exposed the fragility of the existing regulatory framework: ESG declarations, often based on internal audits and third-party certifications, proved vulnerable to sophisticated manipulation.

In response, the **United Nations special rapporteur** has called for a measure that has raised an international clamor: criminalizing the disinformation and greenwashing of the fossil fuel industry. Her position holds that environmental deception is not merely an administrative infraction, but a direct threat to the public trust and the political will needed for climate action [2].

This legal response, however, still lags behind the pace at which new forms of manipulation are generated. Companies, now more aware of the legal risks, turn to artificial intelligence to create "green" narratives that look credible and are, at the same time, hard to dismantle without robust proof. Algorithms can generate misleading ESG reports in a matter of minutes, and social platforms amplify these messages with the help of automated accounts imitating authentic voices.

So while the law begins to react with sanctions and legislative proposals, the phenomenon of greenwashing evolves into more sophisticated, diffuse forms. The need for verifiable mechanisms — like the proof of existence developed in Part IV — becomes urgent: only a layer of trust based on unalterable data can counter the wave of environmental disinformation which, like the smoke of the coal plants, hides the truth and paralyzes the governance needed for a just energy transition.

In this scenario, the fight against greenwashing is not just a matter of fines; it is a battle to preserve the credibility of climate discourse and to guarantee that public policy rests on verifiable facts, not deceptive promises. The legal response must be accompanied by technological tools that ensure the traceability and truthfulness of every environmental claim, because only then can we restore collective trust and advance toward a society grounded in shared truth.

Deepfakes industrialize the lie

In the past two years, generative AI has transformed what used to be an academic curiosity into an instrument for the mass production of deception. The number of deepfakes created worldwide has exploded from **~500,000** in 2023 to **8 million** in 2025, a growth rate that outstrips any recent technological evolution in digital history [3]. This explosion is not a mere numerical increase; it is the manifestation of a new economy of falsehood in which images and audio can be fabricated at small scale with minimal resources, but also at industrial scale by organized entities.

Europol has warned that, if this trend continues, **up to 90% of online content could be synthetic by 2026** [3]. That projection is not just an alarming figure; it is a clear sign that default trust — the assumption that what we see on the internet reflects reality — is reaching its breaking point. When any photograph, video or statement can be artificially generated with a quality almost indistinguishable from authentic material, the thread that holds social and political relations together disintegrates.

The consequences of this phenomenon are multiple and interdependent. First, the **credibility of institutions** is at stake. A manipulated official statement can spread as irrefutable proof of an error or of bad faith, feeding citizen cynicism and sapping the authority of governments. Second, financial markets become vulnerable to **disinformation campaigns** that can move share prices with fake videos of executives or falsified reports of corporate results. Finally, **national security** is exposed: deepfakes can be used to fabricate false statements of state, provoke diplomatic crises or incite hatred and conflict.

This new era of industrial lying is not a mere extension of traditional greenwashing. While the latter relies on manipulating environmental data to improve a corporate image — as the Shein, TotalEnergies and DWS scandals show — deepfakes operate at a more visceral level: they alter the direct perception of what is real. This creates a **cognitive barrier** that prevents citizens from distinguishing reliable information from deception, deepening polarization and shrinking the space for constructive dialogue.

Fighting this threat requires a double strategy: **technological and regulatory**. On the technological side, solutions that anchor the origin and integrity of digital content — like the "t2p" framework we will develop in Part IV — offer a trail for verifying the provenance of what we see. On the regulatory side, legal frameworks must be established to penalize the malicious creation and dissemination of deepfakes, just as the greenwashing sanctioned in 2025 has begun to be [2].

In sum, the industrialization of lies through deepfakes represents a **critical point** for collective trust. If decisive measures are not taken now, society could enter an era in which shared truth ceases to exist and political decisions rest solely on manipulated perceptions. The challenge is great, but the combination of verifiable technology and rigorous regulation offers a viable path to restoring faith in what we see and hear.

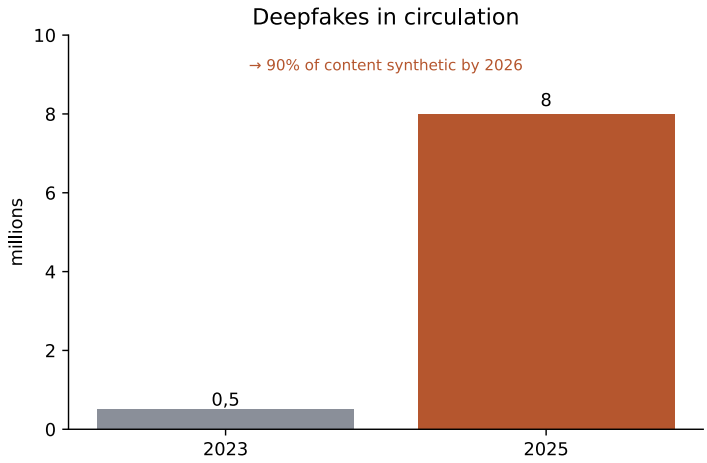


Figure 10.1 — Deepfakes multiply: from ~500,000 (2023) to 8 million (2025), on the way to 90% of online content in 2026.

Why this kills the transition

In an era where climate policies are measured in tenths of a degree and in tons of avoided emissions, trust in data is the invisible fuel that drives decarbonization. Without reliable sources, any target becomes a mirage: investors cannot validate whether a solar plant really produces the energy it announces; regulators do not know whether a petrochemical company's emissions declarations are authentic or doctored; and the public, which must bear the weight of public decisions, faces a labyrinth of uncertain information.

The problem intensifies when the forger does not merely alter numerical data but fabricates entire visual narratives. A video of a wind farm with turbines turning majestically over the coast can come out of an AI with no connection to any real infrastructure. When images and sounds are as easy to fabricate as we have quantified, the traditional criterion of "seeing is believing" evaporates.

This erosion of trust has immediate repercussions in the green capital markets. Greenwashing has reached an inflection point: in 2025 alone, companies like Shein (fined €1.15 million) and TotalEnergies (condemned by a French court for its "carbon neutral" message) were sanctioned, and DWS paid €25 million for exaggerating its ESG credentials [2]. The UN rapporteur has called for such conduct to be criminalized, recognizing that environmental fraud does not only deceive investors; it distorts global climate policy.

The crisis of truth also strikes directly at the targets set in international agreements. A 1998 memorandum of the American Petroleum Institute already advised that "victory will be achieved when average citizens recognize uncertainties in climate science" [1]; today that uncertainty is artificial. If declared emissions are manipulated, progress reports toward carbon neutrality become mere marketing documents, and regulators, lacking tools to verify authenticity, are forced to rely on declarations as false as an AI-generated video.

In short, the inability to identify reliable data paralyzes any credible decarbonization effort. The transition depends on a chain of trust: from the laboratory where CO_2 is measured to the boardroom that approves a renewable energy project. When that chain is broken by disinformation and deepfakes, political decisions rest on manipulated perceptions and society risks being caught in a spiral of doubt and inaction. The solution is not only technical; it requires a rigorous regulatory framework and verifiable technologies that restore faith in what we see and hear.

The promise and the peril

At this moment in history we are witnessing a double explosion of uncertainties: the climate crisis, already measured in degrees and in millions of avoidable deaths, and an accelerating erosion of shared truth that turns every piece of data into a potential deception. The **promise** is clear and unprecedented: for the first time we have tools capable of detecting, unmasking and validating information at global scale, opening the way for society to recover confidence in its own perceptions.

Emerging technologies — manipulation-detection algorithms, content-origin certification systems and open auditing platforms that integrate cryptographic proofs — are beginning to offer ways to trace every piece of information back to its original source, creating a layer of transparency that was previously unthinkable. These solutions become the pillars of the future **t2p**, the framework Part IV will develop to counter the wave of disinformation.

Yet the **peril** lurking behind this opportunity is as real as the climate uncertainties we are trying to mitigate. If we do not act with urgency and regulatory rigor, the crisis of truth will feed a negative loop with the climate crisis: an increasingly distrustful population will refuse to accept scientific data, delaying essential policies such as the decarbonization of the energy sector. The figures we have seen — eight million deepfakes and an internet on the way to being mostly synthetic — sketch a saturation of falsehood in which political decisions are taken on manipulated foundations, leaving society trapped in a spiral of doubt and inaction.

Moreover, the phenomenon of **greenwashing** has become a genuine reputational and legal risk: companies like Shein have paid **€1.15 million** in penalties for misleading sustainability claims [2]; TotalEnergies has been condemned by a French court for its "carbon neutral" message, and DWS has been fined **€25 million** for exaggerating its ESG credentials. The UN special rapporteur has already called for the criminalization of these practices, underlining the risk of ineffective regulation if the technological dimension of the problem is not recognized.

In sum, the **promise** lies in the capacity of the new tools to restore trust and establish a new social contract based on verifiable proof; the **peril**, if we let disinformation keep proliferating without a coordinated response, is that the climate crisis will be hindered even further, leading to a global paralysis in which neither policies nor individual actions can move forward. The key decision stands before us: invest now in verification infrastructure and a robust legal framework, or watch the two crises reinforce each other to the point of irreversible deadlock.

Beyond stepping away from visible falsehoods, we must also confront what Beth Noveck calls the "liar's dividend": the capability of deepfakes not only creates deception, but offers an easy defense to anyone wishing to erase inconvenient data: dismissing it as mere "AI fabrication" [4]. This scenario turns the fake into an instrument of structural denial; any scientific report warning about global warming could be reduced to a digital invention and stripped of weight before the public. Breaking this circle requires a

neutral anchor — an immutable cryptographic record proving that a real fact existed before it was manipulated — beyond detection systems or watermarks like C2PA. Blockchain-based proof of existence, independent of the creator and reliant on no third party, offers exactly that endpoint: it does not try to detect the fake — it certifies the authenticity of the original, and prevents the truth from being dismissed as a mere illusion. The cloned Joe Biden robocall of January 2024, sanctioned by the FCC with an alert against automated calls, illustrates how a neutral proof of existence would have made it possible to check its origin and curb its uncontrolled spread [5].

The double diagnosis

We have diagnosed the double problem: energy and truth. The climate crisis is no longer only a matter of energy sources or decarbonization policies; it now runs parallel to a collapse of shared truth, fed by fossil disinformation, greenwashing and deepfakes. This erosion of public trust paralyzes the governance the energy transition needs in order to advance. Without a common base of reliable information, collective decisions become a game of chance in which manipulated narratives outshine real data.

The signs are undeniable: the American Petroleum Institute's 1998 memorandum already advised that "victory will be achieved when average citizens recognize uncertainties in climate science" [1]. Today, those uncertainties have become a deliberate flood of false information. Greenwashing has gone from a mere marketing device to a legal threat: in 2025, Shein was fined €1.15 million for deceptive practices; TotalEnergies was condemned by a French court; and DWS paid €25 million for exaggerating its ESG credentials [2]. The UN rapporteur is already calling for this behavior to be criminalized, recognizing its devastating impact on market confidence and political will.

In parallel, the deepfake factory has shot up to the volumes that opened this chapter [3]. This reality does not merely disorient public opinion: it compromises national security, judicial processes and international negotiations.

The response must be as double as the problem: we must invest now in robust verification infrastructure and establish legal frameworks that penalize deliberate manipulation. This is where Part IV of the book offers a concrete way out. Artificial intelligence, combined with blockchain, can create an immutable record of origin and authenticity, a proof of existence (t2p) that will allow information to be validated in real time and deepfakes to be confronted with verifiable evidence.

In this context, our call to action is clear: **act now**. Companies, governments and civil society must adopt transparent, auditable certification systems; they must promote legislation that defines severe penalties for greenwashing and the deliberate spreading of falsehoods; and they must invest in advanced media literacy to equip citizens to recognize and reject manipulated content.

As we have seen, we cannot allow the two crises to reinforce each other to the point of irreversible deadlock. Truth is the fuel of trust, and trust is the gear that will turn the wheel of the energy transition. With the right tools — AI, blockchain and determined regulation — we can restore this essential foundation and transform the crisis into an opportunity for a more resilient, transparent future. The moment to act is now; the future depends on our capacity to rebuild shared truth.

---

[1] Union of Concerned Scientists, *Fossil fuel industry's decades of deceit.* https://www.ucs.org/about/news/new-ucs-report-details-fossil-fuel-industry-decades-deceit ↩

[2] United Nations, *Greenwashing*. https://www.un.org/en/climatechange/science/climate-issues/greenwashing ↩

[3] (deepfakes) Blockchain Council / Europol — see research 00 §6. https://www.blockchain-council.org/blockchain/preventing-deepfake-fraud-with-blockchain-content-provenance-verification-workflows/ ↩

[4] Beth Simone Noveck — *Reboot: AI and the Race to Save Democracy* (Yale UP, 2026), ch. 3. ↩

[5] FCC (2024). https://www.fcc.gov/document/fcc-makes-ai-generated-voice-robocalls-illegal ↩

Energy Security Reinvented: Smart Grids and AI

"Safety and certainty in oil lie in variety and variety alone."

— Winston Churchill, House of Commons (1913)

It was a cold October afternoon in 2028 when the core of the Three Mile Island nuclear plant, which for decades had symbolized the worst nuclear accident in North American history, began to hum again. The control monitors showed a steady flow of clean electrons: 837 MW of electricity emerged from Unit 1, rehabilitated by Microsoft and Constellation, feeding a constellation of data centers processing artificial intelligence models for the global management of the power grid.

The project workers, gathered before the great steel doors, watched the power charts rise in a smooth but relentless curve. "This is not just a return to operation," ran the sentiment shared across the control room, "it is proof that the energy security of the future is measured in reliable electrons, not in barrels."

As the turbines spun up, one screen displayed the Microsoft logo beside Helion's, reminding the world that nuclear fusion — still experimental — was the goal of their supply agreement for 2028 [2]. Another screen recalled the 500 MW modular reactor Google had commissioned from Kairos Power, with the first unit scheduled to enter service in 2030 [2]. The same screen announced that DeepMind had cut the cooling energy consumption of Google's data centers by 40%, an efficiency that translates directly into less demand on physical infrastructure and, therefore, more room for renewable sources [1].

Outside, the green fields where fossil fuel storage structures once sprawled were now covered with solar panels and wind turbines which, synchronized with predictive algorithms developed by Google, PJM and Tapestry, balanced the load in real time, avoiding the peaks that had traditionally caused supply failures [1].

The usual noise of a nuclear plant — the constant buzz of the safety systems, the pressure alarms — had been replaced by an almost musical melody: the whisper of the servers optimizing energy distribution and the soft hum of the modernized turbines. The atmosphere was one of renewed confidence; every electron crossing the lines counted as a witness to a new era in which artificial intelligence not only predicted demand but also secured its delivery without interruptions or vulnerabilities.

That moment marked a turning point in the narrative of energy security: it was no longer about controlling oil, but about creating a resilient, intelligent grid capable of sustaining itself and adapting to the needs of the digital world. The reactivation of Three Mile Island was not merely a return to the industrial past; it was confirmation that trust in energy can be rebuilt on digital foundations, where every data point, every AI model and every cryptographic record helps guarantee that the future is fed by clean, reliable electrons.

From oil security to electric security

When Paul Roberts, journalist and energy analyst, addressed energy security in *The End of Oil* (2004), the term "security" came wrapped in a clear image: controlling the flows of oil, protecting the sea lanes and the refinery infrastructure. Today, in 2026, the plot has changed radically; the key is no longer the crude reserve, but the capacity to keep an electric grid clean, abundant and resilient in the face of growing demand driven by artificial intelligence.

This new vulnerability shows itself in several interrelated dimensions. First, the growing reliance on renewable sources — solar and wind — adds a natural variability that makes the grid more sensitive to climatic fluctuations and therefore requires dynamic, real-time management. Second, the proliferation of data centers and AI equipment puts direct pressure on transmission infrastructure: they already accounted for about 2% of the world's electricity in 2022, a figure that could double by 2026 [2]. Third, the emergence of digital layers for certifying energy introduces a cryptographic dimension which, while strengthening traceability, also creates new vectors for cyberattack.

The emerging solutions are as innovative as they are controversial. DeepMind has shown that AI-based optimization can significantly reduce the energy consumption of data centers through better regulation of cooling, and this model is spreading to public distribution networks [1]. Collaborative projects like Google-PJM-Tapestry are modernizing North American infrastructure with predictive algorithms that anticipate demand peaks and redistribute flows in fractions of a second, minimizing unexpected outages [1].

In parallel, the nuclear revival — now in the form of small modular reactors (SMRs) and experimental fusion — offers a stable, high-yield source of electricity. Google and Kairos Power are building a 500 MW SMR with the first unit expected in 2030, as we have seen; Microsoft and Constellation have launched the reactivation of Three Mile Island's Unit 1, with the 837 MW capacity cited earlier [2]. Beyond that, the collaboration between Microsoft and Helion points to nuclear fusion as a target for 2028, promising clean, almost inexhaustible energy that could directly feed the great AI networks without raising emissions [2].

This convergence — AI for grid optimization and advanced nuclear as a constant source — creates an ecosystem in which electric security is measured both in availability and in the capacity to detect and neutralize cyber threats in real time. Every record certifying the origin of a kilowatt is simultaneously a cryptographic proof against manipulation and fraud, while predictive models anticipate potential failures before they occur.

The challenge of 2026 is not to protect a finite resource like oil, but to guarantee a digital-energy infrastructure as adaptable as the AI that consumes it. The vulnerability is no longer the fuel itself, but the network that transports and processes it; only through a synergy of artificial intelligence and advanced nuclear technologies can we build an energy security that is genuinely resilient for the future.

Smart grids and AI

In a Google laboratory in Mountain View, a cold room that once served to keep thousands of servers running now vibrates with the murmur of DeepMind's neural networks. The team cut cooling needs by 40% by optimizing the energy flows of its centers with predictive algorithms that anticipate load peaks and redistribute power toward less congested zones [1]; as we have seen, this success is no academic curiosity but living proof that artificial intelligence can act as the nervous system of an energy network, detecting, processing and reacting to signals in real time with a precision no traditional infrastructure has been able to offer.

This new architecture is taking hold in the United States, where Google, PJM (the independent electricity market operator coordinating the Mid-Atlantic region) and Tapestry have launched a collective project to modernize the national grid. The goal is to integrate a growing share of renewable energy — offshore wind, utility-scale solar photovoltaics — without sacrificing reliability. AI manages flexible demand: it learns from residential and industrial consumption patterns, adjusts tariffs in real time and switches non-critical loads on or off according to the availability of clean energy; thus a factory producing components for electric vehicles can temporarily cut production when the winds are weak, while a community battery park stores the surplus wind energy and discharges it during peak hours.

The distributed-control model rests on a network of IoT sensors gathering data on temperature, current flows, weather conditions and equipment status. Each node processes information locally with learning micro-algorithms to take immediate decisions, while a central AI core optimizes the global strategy. This hybrid approach delivers an almost instantaneous response to unexpected incidents: if a substation is hit by a storm, the AI recalculates the flow routes and redirects electricity through alternative lines within seconds, avoiding massive blackouts.

As we have seen, reducing cooling consumption directly cuts CO_2 emissions and, applied to grid management, multiplies these savings across the entire electric system [1]. Furthermore, the capacity to coordinate intermittent generation and storage in real time reduces dependence on fossil fuels and strengthens energy security: the grids become less vulnerable to external interruptions, since their operation relies on autonomous, adaptive processes rather than rigid infrastructure.

So AI does not merely monitor; it acts as a distributed brain that learns constantly from the system's behavior and adjusts itself to maximize efficiency and resilience. The combination of massive data, advanced algorithms and converging digital infrastructure redefines energy security: it is no longer only about guaranteeing the supply of a scarce resource, but about orchestrating an intelligent network able to sustain itself against the challenges of the future. Roberts's vision of energy as vulnerability has evolved; the key now is to build that synergy between AI and clean energy to secure a world where trust in supply is as robust as the networks that sustain it [1].

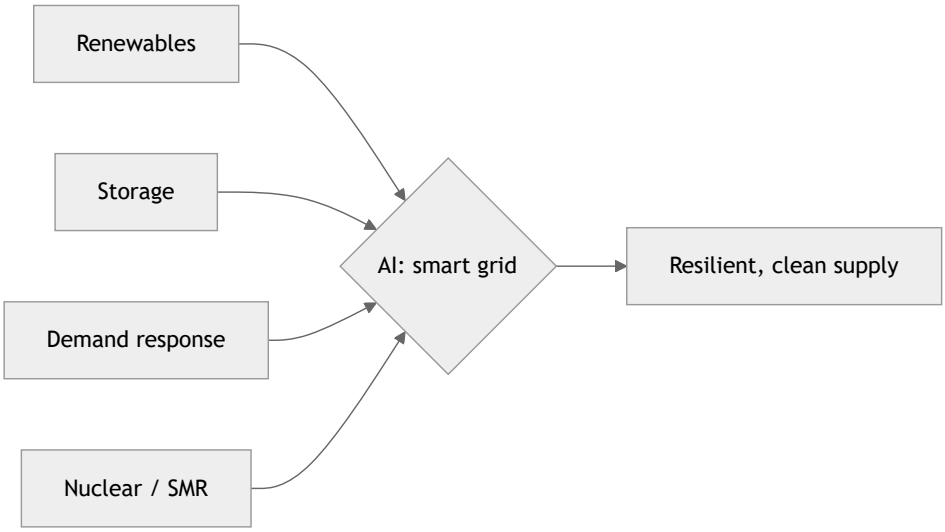


Figure 11.1 — AI orchestrates renewables, storage, demand response and new nuclear into a resilient grid.

The (controversial) return of nuclear

In October 2024 Google commissioned from Kairos Power a fleet of small modular reactors totaling 500 MW, with the first unit expected to operate before 2030; it is proof that AI does not merely consume energy, but steers it toward clean, constant sources [2].

Microsoft and Constellation will reactivate Unit 1 of Three Mile Island, which originally supplied 837 MW, with advanced control technology to generate electricity again in 2028, showing that the stability of smart grids depends on an uninterrupted, carbon-free source [2].

In Ontario, OPG obtained approval in April 2025 to build a GE Hitachi BWRX-300 reactor, with entry into operation expected in 2029 — one of the first approvals of its kind in the West [2]. In parallel, Microsoft and Helion signed an agreement for a fusion reactor to begin contributing energy to the grid around 2028, feeding data centers dedicated to training large AI models and offering a fully clean, almost unlimited energy cycle [2].

These initiatives, however controversial, reveal a clear logic. AI, with its capacity to predict and balance energy demands in fractions of a second, needs an infrastructure that can guarantee constant, emission-free production. SMRs offer modularity, intrinsic safety and shorter installation times than traditional reactors; fusion promises an almost inexhaustible source with minimal waste. In this context, smart grids become the soul of the new energy architecture: algorithms that anticipate demand peaks, redistribute loads and optimize reactor performance to minimize losses and maximize operational efficiency.

The bet, however, is not free of risks. Concentrating critical infrastructure in the hands of a limited number of technology companies can create new points vulnerable to cyber threats or political pressure. Public acceptance of nuclear power, moreover, remains a cultural and regulatory obstacle requiring transparent communication grounded in verified data. The promise of "firm, carbon-free energy at any price" must be accompanied by responsible management of nuclear waste, rigorous safety protocols and a regulatory framework that secures global trust.

In the end, the return of nuclear presents itself as a pragmatic answer to AI's energy needs. It is a bet that combines the visionary spirit of Don Tapscott, co-author of *Blockchain Revolution* — a solution anchored in data and real cases — with an awareness of the inherent dangers, opening a path toward electric grids that are more resilient, more intelligent and cleaner [2].

The promise and the peril

Imagine a power grid that learns in real time from every demand peak, optimizes the distribution of energy the way a doctor adjusts the dose of a medicine and, at the same time, drastically reduces carbon emissions. This is the promise of AI-powered smart grids: an energy security based not on control of the oil routes, but on the predictive and adaptive capacity of autonomous systems. DeepMind has shown that

algorithmic optimization can cut the energy consumption of data centers by up to 40% through more efficient cooling, an improvement that carries straight over to the electric sector when the same principles are applied to managing the national grid [1]; projects like Google-PJM-Tapestry modernize North American infrastructure with AI-based platforms that anticipate failures, redistribute loads and minimize losses, increasing resilience against natural disasters or cyberattacks [1].

In parallel, the rise of small modular nuclear plants (SMRs) offers a constant, clean source of electricity that can feed these smart grids without depending on oil. Google and Kairos Power have already announced a 500 MW SMR with the first unit expected in 2030, while Microsoft and Constellation are reactivating Three Mile Island's Unit 1, with a capacity of 837 MW, for 2028; Microsoft is also working with Helion on a fusion agreement aimed at the same year [2]. As we have seen, the energy reductions reach 40% and the SMR projects foresee 500 MW installations by 2030 and 837 MW by 2028. The combination of reliable nuclear generation and AI-based management creates an ecosystem where energy security becomes an automatic, transparent, measurable function.

Yet the danger hiding behind this vision is just as real. When critical decisions about electricity flows come to be determined by algorithmic models, the integrity of the data and the verifiability of the code become single points of failure; an attack that manipulates the sensors or corrupts the model's training could trigger an artificial overload, a large-scale power cut or even the deliberate diversion of nuclear production into unsafe scenarios. That is why every layer of the infrastructure must be protected by decentralized trust mechanisms — such as immutable ledgers recording sensor values and algorithmic decisions — guaranteeing that any alteration is instantly detected and rejected.

This double face, promise and peril, demands an integrated architecture of trust: AI for optimization and blockchain for verification. Only then can we ensure that the smart grid does not become a new attack vector, but the most robust pillar of the energy security of the future. The key is to develop regulatory frameworks and open standards that mandate the transparency of the models, independent control of their outputs and continuous auditing by independent entities. In this balance between innovation and vigilance, energy security reinvented can leave oil behind and advance toward an era where trust is as secure as the very energy that powers our societies.

Who watches the electrons?

A smart grid does not merely move electrons from one substation to another; it is a stream of data recording every operation, every commercial transaction and every contractual right in real time. This double circulation creates immense opportunities but also exposes vulnerabilities: if the information guiding the algorithms is not reliable, the optimization of generation and consumption can become a source of instability, with consequences ranging from frequency imbalance to massive blackouts. The central question thus becomes who guarantees that these flows are truthful.

The answer emerges as a distributed layer of trust: blockchain. At its core, this technology records every action — from the dispatch of a renewable energy contract to the payment of an end user's consumption tariff — in unalterable blocks verifiable by any authorized participant. This prevents malicious manipulation and enables continuous independent audits. When an AI decides to divert energy toward a data center where the next deep-learning model is being trained, the transaction is recorded and can be reviewed by regulators or external auditors, ensuring there are no hidden deviations or improper gains.

Recent experience shows how this synergy can work at scale. DeepMind has managed to cut the cooling energy consumption of Google's data centers — as we have seen — and if those savings are integrated into a blockchain recording every kilowatt-hour, the transparency of the environmental benefit will be guaranteed [1]. In parallel, the Google-PJM-Tapestry consortium is modernizing the United States grid with AI platforms that balance distributed generation; encoding the flow data and the purchase agreements on a public blockchain would help avoid conflicts of interest and secure the traceability of energy trading [1].

At nuclear scale, projects like the 500 MW SMR developed by Google and Kairos Power (expected for 2030) could incorporate blockchain records for the management of permits, maintenance and operational performance, creating an audit trail consultable in real time by independent regulators [2]. Similarly, the reactivation of Three Mile Island's Unit 1 by Microsoft and Constellation — 837 MW, expected for 2028 — could be documented with a distributed ledger capturing every phase of the safety process and every human or automatic intervention. This approach would be reinforced by the Helion-Microsoft

collaboration on nuclear fusion, aimed at 2028, where energy production data and operating costs could be certified on a blockchain, preventing any manipulation that might compromise market confidence [2].

These examples show how blockchain can act as the backbone of trust in an energy ecosystem that grows ever more intelligent and autonomous. This is not just a technical solution but an institutional framework combining transparency, accountability and continuous audit; the control of oil recedes into the past, and energy security comes to depend on guaranteeing that every electron, every data point and every right moves in a genuinely trustworthy environment. The task for companies, regulators and society is clear: adopt open blockchain platforms, establish independent audit standards and integrate these layers of trust with the AI solutions already transforming grid management, thereby securing an energy security that is resilient, intelligent and clean.

---

[1] Google DeepMind / Data Center Frontier, *AI for the grid.* https://www.datacenterfrontier.com/energy/article/55282719/googles-ai-powered-grid-revolution-how-data-centers-are-reshaping-the-us-power-landscape ↩

[2] IAEA / NucNet, *advanced nuclear and SMRs for AI demand (2024-2025).* https://www.iaea.org/bulletin/data-centres-artificial-intelligence-and-cryptocurrencies-eye-advanced-nuclear-to-meet-growing-power-needs ↩

The (No Longer So) Invisible Hand: Blockchain as a Trust Layer

"What is needed is an electronic payment system based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party."

— Satoshi Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System (2008)

At dusk, the Sant Antoni del Parc neighborhood turns into a small energy exchange. The lights on the facades glow with the soft warmth of the solar panels installed on the terraces and rooftops of the social housing blocks. Each home, equipped with a smart inverter, measures in real time how much energy it generates and how much it consumes. When the sun goes down and production exceeds indoor demand, a notification appears in the Powerledger app on the phone of Marta, who lives on the third floor; with a single tap she can see how many kilowatt-hours (kWh) she has available to sell.

With no intermediary, no hidden bills or additional fees, Marta decides to send the surplus to her neighbor, Joan, who works nights and needs energy for his electronics workshop. Joan receives an alert: "Energy available — accept?" He accepts with a click, and the contract is instantly encoded as a transaction on the Powerledger blockchain, where each unit of energy is represented by a tokenized renewable energy certificate (REC) — one REC equals 1 MWh — [1]. Within seconds, the agreement is immutably recorded: the sensor on Marta's solar panel sends the production data to the blockchain node; Joan's consumption sensors verify that the energy has been received and used; the entire process is logged transparently yet securely.

In the central square, a small monitor shows the neighborhood's transactions in real time: how many RECs have circulated, what percentage of energy has been exchanged among neighbors and what impact it has had on local carbon emissions. The residents watch the chart climb, proud to see their community reducing its dependence on the big energy companies while fostering a circular economy built on green power.

This scene is a perfect illustration of the market's "invisible hand" — this time not a mysterious balancing of interests operating behind the doors of a boardroom, but a visible, auditable network, written in open source. The blockchain acts as a digital witness, ensuring that every transaction can be verified by any participant, with no need for blind trust in a third party. It stands in contrast to the traditional model in which the distribution companies keep the data in proprietary, opaque systems, often inaccessible to consumers.

Still, this experiment remains at the proof-of-concept stage. Most of the 39 projects analyzed show that the potential is real, but they face obstacles such as fragmented registries, inconsistent MRV (monitoring, reporting and verification) data and a lack of legal recognition for tokenized RECs [2]. Even so, the energy efficiency of the underlying layer has improved drastically: Ethereum's Merge is now sustainable — proof that blockchain is no longer a source of waste but a sustainable tool for green energy [3].

In this neighborhood, the invisible hand has become visible. Every click, every token and every entry in the digital ledger turns trust into a measurable asset, opening the doors to a new energy paradigm based on P2P trading, tokenized renewable certificates and near-real-time MRV. The question that remains: when this trust layer is standardized and legally recognized, how many other neighborhoods will be able to replicate the model and completely transform the way we understand and consume energy?

The markets Roberts imagined

When Paul Roberts, in *The End of Oil* (2004), described the energy markets, it was as a great stock exchange: prices set by the invisible play of supply and demand, and opportunities distributed among investors and consumers. His account was not only enthusiastic; it also showed the shadows stretching behind the curtain, with the emblematic example of Enron — a trap of accounting tricks and price manipulation — illustrating the vulnerability of a system where information was the privilege of a few and transparency a luxury.

Roberts called this "opacity" the true enemy of the free market. Transactions were made through hidden contracts, production and consumption data were manipulated by corporate interests, and regulators, often overwhelmed, could not guarantee the truthfulness of the reports. Today, more than a decade later,

blockchain offers a technological answer to that call: its distributed architecture creates an immutable record in which every energy transaction — from the moment the sun generates a kilowatt-hour to the point where it is consumed in the home — is permanently recorded and verifiable by any participant, with no intermediaries able to edit or suppress data.

This new layer of trust is already being tested in the field. Powerledger has developed a platform where neighbors in a district can buy energy directly from one another, bypassing the traditional utility. Every kWh traded is tied to a tokenized renewable certificate; one REC represents exactly one megawatt-hour and is recorded on the blockchain, guaranteeing that the buyer receives verifiable certification of the renewable character of the energy purchased [1]. Advanced sensors installed in the photovoltaic and wind installations transmit production data to the chain, enabling near-instant MRV (monitoring, reporting and verification), and any mismatch between declared and actual production is detected immediately, narrowing the room for fraud.

As the recent survey covering 39 projects indicates, most of these efforts are still proofs of concept [2]; the obstacles are varied: registries fragmented across different platforms, inconsistent MRV data and a lack of legal recognition that prevents the tokens from being accepted as regulated financial instruments. Common standards must be established so that transactions can interconnect at global scale.

At this historic moment, blockchain benefits from a crucial energy evolution: Ethereum's Merge is now sustainable, turning a technology once criticized for its inefficiency into an almost energy-neutral infrastructure [3] — a market where trust no longer depends on an invisible hand, but on a ledger visible to all.

The trust layer

Beyond the neighborhood scene, it is worth unpacking how blockchain acts as a trust layer in three concrete applications. The first is the P2P energy market: protocols like Powerledger enable direct electricity trading between neighbors, with every transaction immutable and verifiable [1]. The key is that the chain removes the traditional intermediaries: trust no longer rests on an "invisible hand" of centralized operators, but on a transparent ledger to which all participants have simultaneous access.

A second application reinforcing this new paradigm is tokenized renewable energy certificates (RECs). Each certificate represents exactly 1 MWh of clean energy produced and is encoded as a unique token on the blockchain, ensuring it cannot be duplicated or forged [1]. This traceability, saturated with transparency, lets consumers and companies check the origin of their energy with a simple scan, turning greenwashing into an unviable practice. Tokenization also simplifies the settlement of certificates: transactions are processed in seconds, without the bureaucracy or costs of traditional registry systems.

Finally, carbon monitoring, reporting and verification (MRV) has been transformed by IoT sensors that write data directly onto the chain. This near-real-time flow of information closes the gaps between physical generation and digital record, significantly reducing human error and data manipulation [1]. With every sensor reading encoded as an immutable transaction, carbon offset projects can demonstrate their real contribution to decarbonization with verifiable evidence.

The energy myth that held blockchain technologies to be a burden on global consumption is now obsolete. After Ethereum's Merge, the network's energy consumption plummeted compared with the previous proof-of-work model [3]; as we have seen, this makes blockchain a sustainable infrastructure for the new energy ecosystem.

Despite these promises, the current landscape is still experimental. A recent survey of 39 blockchain-based projects reveals that most are at the proof-of-concept stage, facing obstacles such as fragmented registries, inconsistent MRV data and a lack of legal recognition [2]. For this trust layer to deploy at scale, common standards and regulatory frameworks guaranteeing interoperability and legal validity must be established. Only then can blockchain consolidate itself as the reliable foundation of a decentralized, transparent and genuinely sustainable energy market.

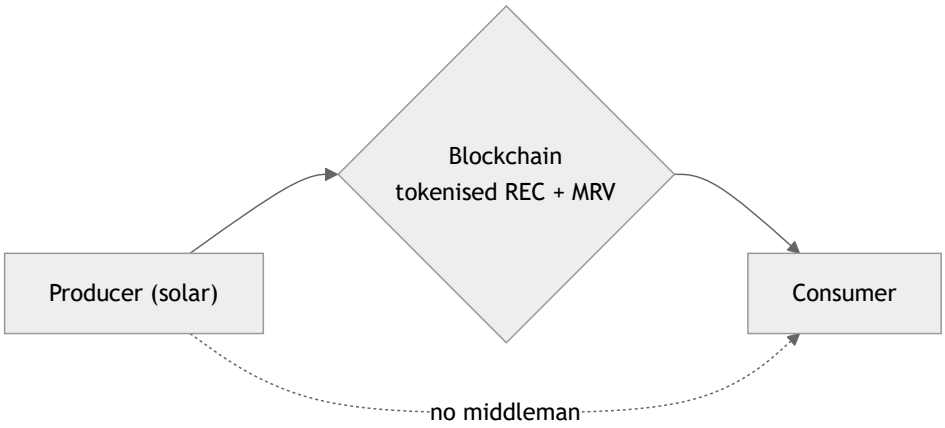


Figure 12.1 — Producer and consumer exchange energy and certificates directly; the blockchain provides trust without an intermediary.

Realism: it is still early

A recent survey tracking the progress of thirty-nine blockchain initiatives devoted to energy finds that most are still at the proof-of-concept stage. These data reveal a landscape where innovation advances with determination but has not yet overcome the fundamental barriers to mass adoption. The obstacles identified are three: fragmented registries that hamper interoperability between platforms; inconsistent monitoring, reporting and verification (MRV) data that compromise the system's reliability; and, finally, the absence of a recognized legal framework legitimizing blockchain-based transactions [2].

This situation should not be read as a technological failure, but as an inevitable stage in the adoption cycle of so disruptive an infrastructure. The promise blockchain offers the energy sector is clear: to create an immutable trust layer guaranteeing that every tokenized renewable certificate (REC) really represents a megawatt-hour produced, and that P2P energy exchanges — like those facilitated by Powerledger — can be recorded transparently and verifiably in near real time [1]. Furthermore, integrating IoT sensors with smart contracts enables almost instantaneous MRV, cutting the latencies that have traditionally weakened carbon certification systems.

Yet the lack of common standards turns this promise into a partial illusion. Each project develops its own tokenization scheme and its own data architecture, generating an ecosystem of "blockchain islands" that cannot talk to each other without costly interfacing. Without clear rules establishing minimum criteria for MRV data quality and the legal validity of the tokens, transactions remain vulnerable to disputes and potential manipulation.

Ethereum's Merge — the upgrade that made the network sustainable — demonstrates that blockchain technology can evolve toward sustainable efficiency [3]. But this structural improvement does not solve the problems of interoperability and legal recognition that still persist. The community must translate this energy advance into a regulatory framework that establishes: (1) format standards for tokenized certificates, guaranteeing that 1 REC really equals 1 MWh regardless of provider; (2) uniform MRV data-exchange protocols, with independent audits and quality certifications; and (3) legal recognition of smart contracts as binding instruments in key jurisdictions.

In this context, close collaboration among regulators, technology companies and energy-sector actors becomes imperative. The authorities can drive the establishment of open certification schemes, similar to the standards that order other industries, while companies like Powerledger must take an active part in defining common APIs that allow interoperability across diverse blockchain networks. Only with this synergy can the promise of immutable trust be turned into an operational, scalable reality.

In short, the current landscape is a delicate balance: the technology has proved its capacity to create a robust trust layer, but it still lacks maturity. Standards and regulatory frameworks are the pillars that must be built now, so that blockchain can consolidate itself as the reliable foundation of a decentralized, transparent and genuinely sustainable energy market. The decision to set these mechanisms in motion is not merely technical; it is a collective commitment to ensure that the new "invisible hand" is truly visible and trustworthy for every actor in the system.

The promise and the peril

Imagine an energy market where every kWh traded carries a small imprint of truth: who produced it, what source generated it, what carbon impact it has and at exactly what moment it was delivered. This is the promise blockchain offers as a trust layer for energy and carbon-credit markets. Instead of blind trust in intermediaries or centralized certifications, transactions are recorded in an immutable ledger where any actor — a community solar generator, a local distributor, an industrial consumer or a regulator — can audit the whole chain in a few clicks, removing the opacity that has fed frauds of fake renewable certificates and failures in emissions-reduction reporting.

Powerledger already enables P2P energy trading between neighbors, writing every transaction into the distributed ledger and linking it to sensors that measure actual production in near real time [1]. Renewable certificates (RECs) are also being tokenized: one REC equals 1 MWh of clean energy and can be transferred like any other digital asset, guaranteeing the traceability of origin and preventing duplication [1]; IoT sensors feed the chain directly, creating a near-instant MRV system that reduces human error and increases confidence in the reported data. As we have seen, a recent survey of 39 energy-related blockchain projects indicates that most still operate as pilots, with fragmented registries, inconsistent MRV data and an absence of regulatory frameworks that blocks interoperability and generates confusion rather than clarity [2].

Furthermore, the trust a blockchain offers depends on the data fed into it. The phrase "garbage in, garbage on-chain" sums up the danger of feeding the blockchain unverified or manipulated information; if the sensors are vulnerable to cyberattack or the certificates are created without rigorous auditing, the chain will immortalize errors and frauds that are hard to correct. That is why establishing quality standards for data sources and robust authentication protocols is indispensable before deploying solutions at scale [2].

Ethereum's Merge has defused the energy concern that traditionally surrounded blockchain technologies, making the protocol sustainable and turning the infrastructure into a viable option for high-volume applications like energy markets [3]. This removes the apparent obstacle of "hidden costs" and allows the focus to rest on the key question: guaranteeing that the data are reliable, verifiable and compatible across every part of the ecosystem.

In short, blockchain can be the trust layer that transforms energy and carbon markets into transparent, fraud-resistant systems — provided common standards are developed and the origin of the data is protected. Without these foundations, the revolutionary potential could remain trapped in a web of incompatible silos, leaving the "invisible hand" as untouchable as before.

The social contract of energy

In an ecosystem where every kilowatt-hour can be bought, sold or exchanged between neighbors, trust must be verifiable and immutable. Here enters the central concept of this book, **t2p**, the blockchain-based proof of existence that closes the circle of trust: a cryptographic mechanism guaranteeing that certified energy or carbon credit cannot be altered without leaving an inescapable trace.

Imagine a community of self-consumers in Barcelona that generates solar energy and sells it on a P2P market. Each generator issues a **tokenized REC** (1 REC = 1 MWh) representing the right to sell that verified energy [1]. When it is recorded on the blockchain, a proof of origin is included: the smart meter's data, the exact date and time and the renewable energy provider's certification. The blockchain engraves this information in a distributed ledger that any actor can consult without doubting its authenticity.

This proof-of-existence layer becomes the foundation for near-real-time **MRV (monitoring, reporting, verification)**. Sensors installed along the network collect production and consumption data; the data are cryptographically summarized and added to blocks on an infrastructure that, after **Ethereum's Merge**, keeps its energy consumption minimal [3]. Every transaction is transparent and sustainable.

Even so, the road to mass adoption is still strewn with obstacles. A recent survey of 39 projects indicates that most remain at the proof-of-concept stage, limited by **fragmented registries** and inconsistent MRV data [2]. Without a common regulatory framework, the authorities do not legally recognize these tokens, and market actors are reluctant to rely on an infrastructure not yet fully standardized, generating silos of trust where each platform runs its own protocol.

The solution lies in establishing **open standards** defining how a proof of existence must be issued, signed and validated. Standards organizations can lead this process, ensuring the proofs are legally recognized across jurisdictions; once these standards are in force, blockchain will cease to be an experimental technology and become the invisible yet palpable layer of trust enabling P2P markets, tokenized renewable certificates and transparent MRV.

In short, **t2p** is not just a technical tool; it is the social contract of the energy future. It ensures that what enters the chain is authentic, creates an immutable record of trust and offers a solid basis for the climate and economic policies that depend on reliable verification. The next step is clear: adopt global standards, integrate proofs of existence into every system and let blockchain become the true core of trust of the new energy system.

---

[1] RWA.io / Powerledger, *tokenized RECs and MRV on blockchain.* https://www.rwa.io/post/renewable-energy-credits-on-blockchain-a-tokenized-approach ↩

[2] Frontiers in Blockchain (2025), *state of blockchain carbon markets.* https://www.frontiersin.org/journals/blockchain/articles/10.3389/fbloc.2025.1603695/full ↩

[3] ethereum.org, *The Merge*. https://ethereum.org/roadmap/merge/ ↩

t2p: The Proof of Existence Against Deepfakes

"The struggle of man against power is the struggle of memory against forgetting."

— Milan Kundera, The Book of Laughter and Forgetting (1979)

March 2026, a sunny afternoon in the Barcelona-Sabadell area: the chief executive of one of Spain's largest electricity companies appears in a video circulating on Twitter. In just 38 seconds, the footage — an elegant office with panoramic views over the Vallès Technology Park — shows the executive, gaze serious and a folder full of emissions charts in hand, declaring that "we have decided to adjust the emissions data to better reflect operational reality." The sentence, delivered in an apparently transparent tone, instantly becomes the stage for an unexpected crisis.

In less than an hour, the company's share price falls 12% on the Iberian exchange. The news multiplies: financial analysts warn that investors now doubt every ESG certification and fear a possible sanction from the market regulator. The media, from La Vanguardia to Bloomberg, run alarming headlines — "The truth erased: the great deception of the energy transition." At this moment, social networks fill with cybersecurity experts recalling Europol's latest reports: deepfakes have gone from some 500,000 in 2023 to 8 million in 2025 [2].

Seven hours later, a digital forensics firm publishes a detailed report: the video is a deepfake. Through advanced detection algorithms and comparison with original metadata, they demonstrate that the executive's face was generated by a neural network trained on public images and that the audio was synthesized from voice samples available in corporate archives. The doubt, however, no longer goes away. Investors, regulators and the public find themselves trapped between the certainty of an official statement and the uncertainty of a technological deception.

This scenario illustrates an increasingly palpable reality: the transition toward a decarbonized economy depends on reliable data. ESG reports, carbon MRV (Measurement, Reporting and Verification) methods and clean energy certifications are the foundation on which multibillion investments and global climate policies are built. When visual truth can be altered in a matter of minutes, these layers of information lose credibility and, with them, the trust needed to drive large-scale renewable energy projects.

In this context, the need for an immutable trust layer becomes evident. Modern blockchain offers exactly this: a decentralized registry where every piece of data — from emissions certificates to smart contracts enforcing energy limits — can be verified without depending on a single point of failure. The proof of existence, called t2p, is born as the flagship answer: a cryptographic mechanism that records the time and the content of a document or media file on a blockchain, creating an indelible fingerprint that anyone can audit.

This scene, as sober as it is cinematic, does more than capture the urgency of the moment; it highlights the chapter's central paradox: without a robust system to guarantee the authenticity of data, the energy transition remains trapped in a fog of uncertainty. The solution lies in the hands — and in the code — of the proof of existence, which promises to restore lost trust and let genuine projects advance without the shadow of deepfakes smothering them.

The stakes

As we have seen, the number of deepfakes has gone from some 500,000 in 2023 to 8 million a year, and Europol projects that by 2026 up to 90% of online audiovisual material could be artificially generated, a reality that transforms the information landscape and, by extension, that of energy data [2].

This wave of digital forgery is met by a regulatory framework that is still nascent but increasingly firm. The **EU AI Act**, adopted in 2024, obliges platforms to clearly label synthetic content and to provide a **watermark** indicating its provenance [3]. In the United States, the **Take It Down Act** establishes penalties for anyone hosting or distributing deepfakes without proper labeling. And since January 2025, **CISA** (the Cybersecurity and Infrastructure Security Agency) has published an advisory urging critical infrastructure companies — including electricity generation and distribution — to strengthen verification of data integrity before using it in ESG reporting [3].

Without a reliable layer of trust, the transition to a decarbonized economy becomes vulnerable to manipulation. Regulators need clean data to validate compliance with climate targets; investors look for

MRV (monitoring, reporting and verification) certificates guaranteeing that announced emissions reductions are real and not mere marketing narratives. The public, increasingly aware, demands transparency: if a photo of a new solar plant can be a deepfake, the project's credibility disintegrates in seconds.

In this context, the **proof of existence** emerges as an indispensable technological answer. In its best-known form — the **OpenTimestamps** protocol — this cryptographic layer records the exact, immutable time of a document or dataset, ensuring that any later attempt at modification is exposed [1]. The integrity of the record is consolidated on the blockchain, where each hash — including perceptually similar versions developed in advanced **perceptual hash** research — is verified and linked to the previous block, creating an unbroken chain of authenticity [5].

Moreover, the evolution of **Ethereum's Merge**, which cut its energy footprint by **99.95%**, shows how blockchain infrastructure can be sustainable and robust for this purpose at the same time [6]. A proof of existence inserted into a green blockchain therefore not only guarantees the truthfulness of the content, but does so without compromising climate goals.

The stakes are clear: **if we cannot trust the data**, no policy or investment for decarbonization can advance with legitimacy. The combination of demanding regulation and a blockchain-based trust layer offers the framework needed for truth to become the anchor of the energy transition. Companies, governments and civil society must adopt this proof-of-existence architecture now; only then can we keep deepfakes from becoming the insurmountable obstacle to a clean, trustworthy future.

What the proof of existence is

Imagine you have before you a decarbonization report with thousands of sensor readings, ESG certificates and carbon MRV (monitoring, reporting and verification) calculations. You want to prove that the document existed on a specific date and that its content has not been altered since. The answer lies in the proof of existence, a three-step mechanism that turns any file — from spreadsheets to turbine inspection videos — into an immutable fingerprint recorded on the public blockchain.

Step 1 — SHA-256 hash of the file. The first step applies the SHA-256 cryptographic algorithm to the file's content. The algorithm produces a 64-character string — a unique, one-way digital fingerprint — representing the entire document. Its key property is irreversibility: any change, even a single byte, generates a completely different fingerprint. This guarantees that the file never has to leave the device to be processed; privacy stays intact because only the small hash string, not the original content, is transmitted.

Step 2 — Merkle tree. Once we have a set of hashes (one file or thousands of related files), they are organized into a hierarchical structure known as a Merkle tree. Each pair of hashes is combined and hashed again, creating intermediate nodes up to the root — the "Merkle root." This root consolidates the entire dataset into a single fingerprint representing the complete integrity of the block of information. The advantage is twofold: any individual element can be verified by checking only its path to the root, and the amount of data that must be transmitted for verification is drastically reduced.

Step 3 — Anchoring on a public blockchain. The final step engraves the Merkle root onto an immutable blockchain, such as Bitcoin, through timestamping services like OpenTimestamps [1]. The process sends the root to the Bitcoin network, which includes it in a mined block and assigns it a verifiable timestamp. Once the block is confirmed, a .ots proof is generated containing the cryptographic path from the Merkle root to the hash of the block where it was inscribed. This proof can be downloaded and stored locally without depending on third parties.

Verification is fully autonomous: any reader recomputes the SHA-256 of the original file, rebuilds the corresponding Merkle tree and compares its root with the one in the .ots proof. If even a single byte has been modified, the comparison fails immediately and demonstrates the manipulation. No central authority needs to be trusted; security rests solely on the mathematical properties of the hash algorithm and the immutability of the blockchain.

This procedure turns the proof of existence into a trust layer that can be integrated into any ESG or MRV workflow. Energy companies can register their emissions-reduction certificates, regulators can verify the authenticity of reports without on-site audits, and civil society gains a tool to validate the truthfulness of the data underpinning the transition to a clean future. In a world where deepfakes threaten the credibility

of any visual or documentary claim, the proof of existence offers a clear, verifiable technological answer — a digital shield against disinformation that protects the very core of the energy transition.

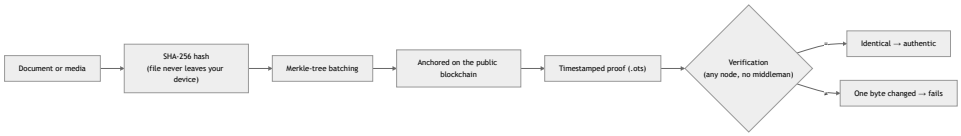


Figure 13.1 — From the hash to decentralized verification: no one has to trust anyone; only the mathematics.

Why blockchain, and why now

The first time the public heard about the energy cost of blockchains, the image that stuck was of a mining farm feeding the Bitcoin network with hundreds of megawatts — a monster devouring electricity as if it were fuel. That scenario made the word "blockchain" synonymous with waste and inefficiency, and the criticism spread to the executive tables that judged the technology an unacceptable expense for clean energy projects.

The reality has changed radically. The move from the proof-of-work (PoW) model to proof-of-stake (PoS) is not merely a technical evolution but a revolution in the balance between security and consumption. The operation known as Ethereum's "Merge," completed at the end of 2022, replaced its PoW architecture with a PoS mechanism that has cut the protocol's energy consumption by **99.95%** [6] — as we have seen; a network that once required energy comparable to the annual consumption of a small country now runs on the equivalent of a mid-sized city. The marginal cost of adding a transaction becomes almost nil, while security is maintained by the economic incentive of staking.

Modern PoS chains, building on this model, are now able to provide a **trust layer** at minimal energy cost and with a scalability that suits them to mass applications. A single existence anchor, recorded in a transaction containing the Merkle root of a collection of documents, can validate thousands — even millions — of files with a single hash. This efficiency translates into a drastic reduction in computing resources and, therefore, a small carbon footprint for the certification process.

The PoS model also offers benefits beyond consumption. **Decentralization** remains robust because participation does not depend on costly specialized hardware, but on tokens that any actor can acquire and "stake." This democratizes access to the verification infrastructure, allowing small entities — local renewable energy cooperatives, environmental monitoring NGOs — to join a global network without investing in large data centers.

In this context, the proof of existence (t2p) becomes a practical, reliable solution to the crisis of truth that threatens the energy transition. The trust a blockchain layer provides is not a "hidden cost" but a minimal investment compared with the economic and social value of an information system that is verifiable, transparent and resistant to manipulation. ESG data, carbon MRV reports and any other record critical to the credibility of decarbonization can be anchored to this infrastructure without fear that the act of certification becomes a source of pollution.

In short, the myth of blockchain's energy waste belongs to the PoW past. Today's PoS chains offer **trust at minimal cost**, scalability and accessibility, making them a viable trust layer for the infrastructure that sustains the transition to a clean future. The moment to adopt them is now: only then can we guarantee that the data driving climate change action are as reliable as they must be.

C2PA and the proof of existence, complementary

Within the framework of the **Content Credentials** driven by the **C2PA** initiative, the first step toward truthfulness is to record the origin of content at the exact moment of its creation. Each file receives a SHA-256 cryptographic fingerprint which, together with a Merkle tree and an X.509 signature from an accredited provider, forms a chain of trust that can be verified without depending on intermediaries [4]. This architecture has proved its worth in sectors like digital journalism and audiovisual production, where platforms can certify that a video or an image comes from a legitimate source and has not been altered afterward.

Yet this model has a structural limitation: **trust depends on the creator's cooperation**. If the provider decides to withdraw the credential — for commercial or legal reasons, or simply through neglect

— the absence of that proof offers no useful information about whether the content has been manipulated or not. In other words, **the deletion of the credential amounts to null evidence**, and the system is left unable to prove that a certification ever existed [4]. This creates a critical vulnerability in an environment where deepfakes are spreading at alarming rates: according to **Europol** data, the number of deepfakes in circulation has gone from some 500,000 in 2023 to 8 million in 2025, with a Europol projection putting up to **90% of online content as synthetic in 2026** [2].

The **proof of existence** (t2p) offers an answer to this gap. Based on the same principle as **OpenTimestamps**, this solution records the content's hash directly on a PoS blockchain, creating an **independent, immutable** anchor that persists even if the original metadata disappear or are erased [1]. The key characteristic is that the proof requires no prior authorization: any user can generate the hash and publish it on the network, ensuring the moment of creation is immortalized. This entirely eliminates the risk of a "silent deletion" and provides an additional layer of trust that transcends the limitations of the C2PA model.

This is not a rival solution but a **synergy**. Combining C2PA's origin certification with t2p's immutable anchoring forms a defense ecosystem on four fronts: (1) verifiable provenance, (2) an irreversible temporal anchor, (3) automatic detection based on perceptual hashes — a line of research that in 2025-26 has successfully linked them to digital media stored on blockchain [5] — and (4) media literacy for end users.

This alliance is especially relevant for the critical infrastructure of the energy transition. ESG reports and carbon MRV mechanisms depend on **reliable traceability** of energy production, transport and consumption data. Without a robust trust layer, any decarbonization effort is undermined by doubts about the truthfulness of the claims. Implementing C2PA and t2p simultaneously means guaranteeing that emissions certificates, reduction reports and audits are **not only auditable, but impossible to alter**.

In conclusion, the convergence of C2PA and the proof of existence closes a narrow technical gap and creates an **integral trust framework** capable of supporting the most demanding climate policies and countering the wave of digital disinformation. The call to action is clear: companies, regulators and developers must adopt both technologies at once, establishing open protocols that integrate content credentials and blockchain anchoring into their platforms. Only then can we guarantee an energy transition built on true data that can withstand the bombardment of deepfakes.

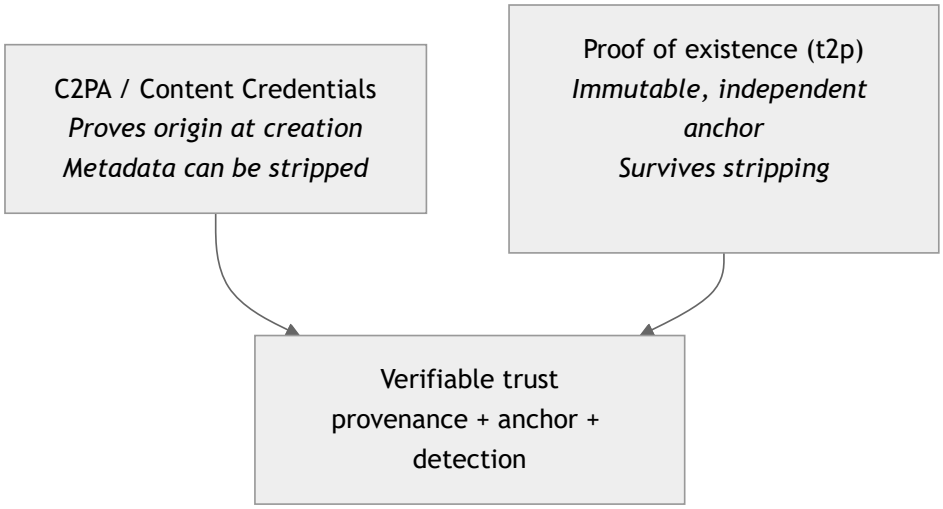


Figure 13.2 — C2PA proves "who and when" at origin; the blockchain anchor proves "this existed and has not been altered."

t2p, the flagship solution

Here is where theory reaches the pocket. t2p turns a phone into a portable proof studio: every photo, video, audio recording, scanned document, web page or screen capture receives, at the instant of its creation, a unique cryptographic fingerprint — a **SHA-512** hash — recorded on a **public blockchain** — today, Polygon's; the app can implement any other Ethereum L2, because the proof does not depend on any particular chain, but on the fact that the anchor is public and verifiable by all. The result is tamper-proof evidence that something existed and when, verifiable by anyone and forever. And one design

principle governs everything: the original file **never leaves the device**; only its fingerprint travels. t2p cannot see what you capture.

The flow is the proof of existence described above — as we have seen — but applied with forensic rigor. Upon capture, the app computes the hash from the exact bytes of the file and anchors it to the public blockchain with a timestamp. It adds, in addition, an **RFC 3161** trusted timestamp (the eIDAS standard) independent of the chain and, if the user allows it, **forensic-grade GPS** coordinates: they are sealed only if the signal is genuine — accuracy of about 50 meters and a maximum age of 30 seconds; a mocked location blocks certification to protect evidentiary value. Each certification generates a professional **PDF certificate** with the hash, the transaction identifier, the RFC 3161 seal, the location and a **QR code** to check it on the chain explorer.

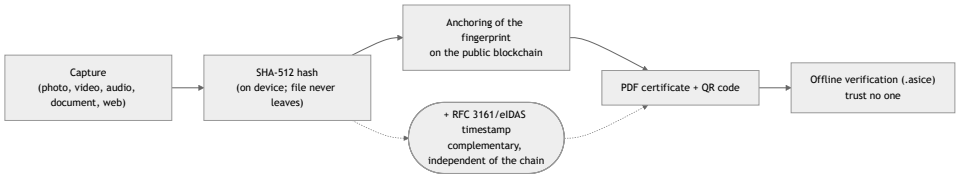


Figure 13.3 — The t2p flow: capture, seal the fingerprint on the public blockchain, and verify it even without a connection or intermediaries.

The most solid piece for legal use is the `.asice` **evidence package**: a standard ASiC-E container (ETSI EN 319 162) that includes a byte-for-byte copy of the file, the certificate, the timestamp and its chain of trust, and even a text file with the `openssl` commands to check it. Anyone can verify the proof **completely offline, without the app and without having to trust t2p**; it is, in essence, the difference between "trust us" and "check it yourself."

On this foundation, the use cases unfold on their own.

- **Journalism and media.** A reporter certifies a photo or a video at the exact moment of capture; the reader verifies its origin by scanning the certificate's QR code.
- **Legal evidence and chain of custody.** The forensic GPS, the RFC 3161 seal and the `.asice` package make a piece of evidence withstand courtroom scrutiny. And the **video-selfie signature** — with liveness checks and biometric verification — ties a document to a real, present person.
- **Intellectual property and authorship.** Sealing a work proves it existed on a certain date, without registries or intermediaries.
- **Integrity of climate and ESG data**, the bridge to this book's thesis. A company can certify an emissions sensor reading, an MRV report or the photograph of an installation in such a way that no one — not even the company itself — can alter them afterward. The transition needs exactly this class of incontrovertible data.

t2p embodies an idea as simple as it is powerful: trust should not have to be requested — it should be checkable. In an age when any image can be fake, making truth verifiable — and doing so from your pocket — is perhaps the most discreet and, at the same time, most decisive infrastructure of the entire transition.

The promise and the peril

The proof of existence (t2p) offers **verifiable trust** that can be demonstrated at marginal cost: each data hash is recorded on the public blockchain, where immutability and transparency are guaranteed by the decentralized network. This transforms a claim — for example, "this ESG report reflects the plant's real emissions" — into cryptographically proven evidence that any actor can consult without needing to trust an intermediary. The marginal cost of adding a proof is minimal: a single anchoring transaction can validate thousands of hashes at once through the Merkle tree. This drastic reduction, the fruit of a fusion of technologies (Ethereum's Merge, −99.95% in energy consumption) and layer-2 solutions, makes the proof of existence **scalable to millions of insertions a day** without creating economic barriers for NGOs or small companies [6].

At the operational level, the benefit is clear: regulators can validate carbon MRV (monitoring, reporting and verification) reports in real time; investors can audit a project's trajectory without relying on centralized certifications that are often subject to political pressure. The **trust layer** provided by the blockchain thus becomes the foundation on which the credibility of the entire energy transition rests.

Without this support, deepfakes — which according to Europol could reach 90% of online content by 2026 — erode public faith and cause any ESG data to be viewed with skepticism [2].

Yet this promise comes with an inherent danger. **De-anonymization** is an inevitable consequence when proofs of existence are linked to identifiable metadata. The C2PA consortium has noted that the disclosure of digital signatures can reveal the origin of content, exposing journalists and activists to reprisals in authoritarian regimes [4]. This demands a careful privacy design: proofs must be kept separate from any personal information and, where necessary, use **perceptual hashing** to protect the content without compromising verifiability [5].

Furthermore, the proof of existence **does not detect fakes by itself**; an attacker can create a fraudulent document, register its hash and present it as genuine. The system's security depends on complementary processes — independent audits, certifications from reliable sources and revocation mechanisms — that must be integrated into the ecosystem. Excessive faith in the technology could lead to **regulatory laxity**, with institutions delegating all control to the chain and neglecting the need for human oversight.

In conclusion, the proof of existence is a **necessary but not sufficient technology** for guaranteeing truth in the age of deepfakes. Its success depends on building an ecosystem that combines blockchain transparency with robust regulatory frameworks (EU AI Act, Take It Down Act, CISA) and privacy principles designed to protect the most vulnerable [3]; only then can we turn the promise of verifiable trust into a sustainable reality for the energy transition.

Moreover, Beth Noveck confirms that decentralized technologies — blockchain among them — are one of the proposed avenues for guaranteeing the authenticity of digital content and increasing its transparency [7]. Yet she argues clearly that watermarks and provenance standards, on their own, do not solve the problem, because they depend on the cooperation of the creator and of the platforms that implement them [8]. This is where the neutral anchoring of **t2p** shows its full potential: recording the hash on a public chain with no intermediary or central authority ensures that the proof of existence remains independent of the creator and of trusted third parties, closing the gap Noveck identifies and reinforcing the promise of verifiable, universal trust.

It's time to act

You will find that truth is no longer a natural condition of public debate; it is a resource that must be bought and protected. That is why each social group has a concrete mission in the **t2p** ecosystem.

1. Citizens: the first filter

Before sending or sharing any video, audio or image, ask for the proof of existence. Tools based on OpenTimestamps make it possible to verify that a file was created before a given moment and has not been altered since [1]. This simple step — a click on a browser extension or a mobile app — can prevent the mass propagation of deepfakes, a wave Europol documents on the rise [2]. Every user who becomes a watchdog is an additional layer of security: trust is built not only at the institutional level, but also through a culture of constructive doubt.

2. The media: the commitment to publish with proof

News outlets must make the **timestamp** a requirement for all visual content that could influence public opinion or political decisions. That means integrating a blockchain verification API into the editorial workflow and adding, as visible metadata, the perceptual hash certified on the chain (see the studies demonstrating these hashes' resistance to subtle manipulations [5]). When an outlet states "verified with t2p," it offers a tangible guarantee that goes beyond declarations of good faith. The absence of these practices has already brought sanctions under the EU AI Act, which provides for fines of up to 7% of global annual turnover for the gravest infringements [3].

3. Companies: anchor ESG and MRV data

Corporations that report carbon emissions or ESG certificates must store their reports on an immutable public chain. The migration to the Ethereum network, which has cut energy consumption almost entirely with the latest proof-of-stake merge [6], offers a sustainable environment for this purpose. Each data record is signed with a private key and associated with a perceptual hash that guarantees the integrity of the original document. Moreover, interoperability between standards like C2PA (which currently shows

limitations in origin authentication) and timestamping solutions makes it possible to create a hybrid trust layer: the blockchain contributes immutability, while the metadata protocols add context and a trail of ownership [4].

4. Regulators: define standards and foster interoperability

The authorities must establish regulatory frameworks that mandate the use of proofs of existence for any information with social or environmental impact. That means:

- **Interoperability requirements** between timestamping solutions, C2PA and the initiatives of the Take It Down Act, to ensure certificates are recognized across the EU and the United States.
- **Independent audits** to validate the implementation of the protocols at scale, with quarterly public reports on the compliance status of companies and media.
- **Tax incentives** for organizations that adopt low-energy blockchain solutions, such as Ethereum's Merge.

5. A future without verifiable truth is not a sustainable future

Without a verifiable layer of trust, the energy transition loses its credibility: investors, citizens and governments will be unable to rely on the data underpinning decarbonization projects. The proof of existence **t2p** offers the mechanism to turn climate promises into verified facts.

The call, then, is clear: **act now** — verify before sharing, publish only with proof and anchor your data on the blockchain. Only then can we advance toward a credible, lasting energy transition. It's *time2act*.

---

[1] OpenTimestamps — *a timestamping proof standard*. https://opentimestamps.org/ ↩

[2] Blockchain Council — *Preventing deepfake fraud with blockchain*. https://www.blockchain-council.org/blockchain/preventing-deepfake-fraud-with-blockchain-content-provenance-verification-workflows/ ↩

[3] London Blockchain — *Can blockchain save truth? (2025)*. https://londonblockchain.net/blog/blockchain-in-action/can-blockchain-save-truth-tackling-deepfakes-and-disinformation-in-2025/ ↩

[4] C2PA — *Content Credentials Explainer (official specification)*. https://spec.c2pa.org/specifications/specifications/2.4/explainer/Explainer.html ↩

[5] arXiv — *Provenance Verification via a Perceptual Hash Registry Anchored on Blockchain*. https://arxiv.org/html/2602.02412v1 ↩

[6] ethereum.org — *The Merge*. https://ethereum.org/roadmap/merge/ ↩

[7] Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 3. ↩

[8] Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 3. ↩

Hydrogen, Capture and the Bets Still Missing

"I believe that water will one day be employed as fuel, that hydrogen and oxygen which constitute it, used singly or together, will furnish an inexhaustible source of heat and light."

— Jules Verne, The Mysterious Island (1874-75)

It was an October afternoon in 2025 in Immingham, in the north of the United Kingdom, where the sea cold seeps between the metal structures of one of the most anticipated industrial complexes on the continent. The steel towers rose like giants promising to transform energy into green hydrogen, an invisible flame of a sustainable future. On paper, the project was to produce hydrogen by electrolysis fed with renewable energy and to become one of the pieces of Europe's industrial decarbonization.

But that morning, the sound that broke the routine was not the quiet hum of the compressors but a sudden announcement: **Air Products had decided to cancel the Immingham green H_2 project**, an initiative valued at some £2 billion and considered one of the biggest global bets on industrial decarbonization. The message, short but blunt — "Project suspended indefinitely" — spread instantly through the media and across the networks where engineers and industry analysts began commenting with a mixture of disbelief and resignation.

The news was not just a financial setback; it reflected a phenomenon that had been consolidating over recent years: **an oversized dream about green hydrogen**. From the initial promise onward, the sector had poured immense resources into promising a clean energy capable of replacing fossil fuels in every sector, from maritime transport to cities. Yet 2025 marked a **brutal reset** for the market. According to consolidated reports, roughly **60 large green hydrogen projects were canceled that year**, wiping out a planned capacity of **−4.9 million tonnes per year** and cutting the projected global pipeline for 2030 from some **49 to just 37 million tonnes** [1].

At the heart of the complex, the towers that had been erected as symbols of a new era fell silent. The engineers who had spent years optimizing the electrolysis processes found themselves looking at their equipment with a gaze wavering between frustration and acceptance: the cost of green hydrogen had risen **25-40% since 2022**, now standing between **€4 and €7 per kilogram** — almost three times more expensive than its gray predecessor, still hovering around **€1.5-2.5/kg**, and clearly above blue hydrogen with carbon capture (CCS), which held to a range of **€2-3.5/kg** [2].

The conclusion echoing through Immingham traveled far beyond the complex's walls; it became an echo running down the global supply chain, reminding every actor in the sector that optimism without data can be as dangerous as inaction. Green hydrogen, as it had been imagined — energy for everything — was not viable at scale at that moment. The awakening from this oversized dream marked the beginning of a new stage: a call for honesty and for a revaluation of the bets, focusing on concrete applications — steel, fertilizers and chemicals — where green hydrogen can genuinely deliver substantial change.

The hydrogen dream

When Paul Roberts, in *The End of Oil* (2004), devoted to hydrogen a chapter ironically titled *And Now for Something Completely Different*, this gas was presented as the keystone of a clean economy: a fuel that, when burned, would leave only a trace of water and could therefore replace oil without a carbon footprint. The vision, so seductive, mobilized investors, governments and large energy conglomerates; impassioned forecasts announced a rapid transition toward a global network of hydrogen stations, turbines fed by this clean gas and a fully sustainable "life cycle."

That hope was fed by spectacular numbers and by an avalanche of industrial-scale projects, such as Air Products' mega-investments in Immingham or BP's project in Duqm, Oman. The idea was clear: if millions of tonnes could be produced annually, the market would adjust and hydrogen would become the new energy base.

But the reality of 2025 broke this tale of optimism with a call to honesty. Roughly sixty large green hydrogen projects were canceled that year and, as we have seen, planned production was slashed while prices rose between 25% and 40%, settling into a band of €4-7/kg [1].

Furthermore, the infrastructure needed to distribute hydrogen at global scale — storage tanks, pipelines and filling stations — still faces technical and regulatory barriers. Losses in transport and the difficulty of

storing it safely make the "universal hydrogen" proposition less attractive than alternatives like direct electrification. In parallel, carbon capture and storage (CCS) advances slowly and remains expensive, leaving blue hydrogen with CCS in an uncertain position [2].

Thus the original promise — a clean fuel that leaves only water — has been diluted by economic and technological limits. Green hydrogen still has a crucial role, but no longer as a universal solution: it is a specialized complement in sectors where carbon is intrinsically hard to eliminate — steel production, fertilizers and basic chemical reactions. In these niches its impact can be measurable and justifies targeted investment, but the "energy for everything" narrative must be abandoned in favor of a perspective grounded in concrete data and well-quantified risks.

The reset of 2025

2025 marked an inflection point that shook the most optimistic expectations for green hydrogen. In a year that opened with promises of massive investment and an avalanche of projects announced as pillars of the energy transition, the market ran into hard data: roughly 60 of the large green H_2 projects were canceled. This retreat represents an estimated loss of 4.9 million tonnes per year of potential production and has cut global green H_2 projections by nearly a quarter from the original 2030 plan. The expectation now stands at 37 million tonnes per year — a drastic fall compared with the 49 million foreseen two years earlier — made concrete in emblematic cases such as Air Products' project in Immingham (United Kingdom), BP's initiative in Duqm (Oman) and ArcelorMittal's effort to decarbonize its steel plants in Germany [1].

This wave of cancellations is not random; it reflects an economic model too fragile for market reality. Many projects had been sized to capture subsidies and government incentives without rigorous analysis of real short-term demand. The weakness of industrial orders — especially in sectors still heavily dependent on fossil carbon — left these facilities with costly infrastructure and too few customers to justify amortizing the initial outlay.

Moreover, the cost of producing green hydrogen has soared between 25% and 40% since 2022; it currently ranges between €4 and €7 per kilogram, far above the price of gray hydrogen and of blue hydrogen with carbon capture — as we have seen. This gap makes green hydrogen competitive only in applications where emissions reduction delivers critical added value — steel production, fertilizers and certain chemical processes — while it remains barely viable as a universal substitute for fossil fuels in sectors like heavy transport or baseload electricity.

The combined impact of insufficient demand and rising costs has generated a negative spiral: investors, spooked by uncertain returns, have cut their contributions, leaving more projects unfunded and reinforcing the perception that green hydrogen is still a risky bet. This scenario forces a rethink of hydrogen's role in the energy mix: not as a universal solution, but as a strategic component for niches where emissions reductions translate directly into competitive and regulatory advantages.

In short, 2025 acted as a brutal "reset" that forced the industry to abandon universalist narratives of the future and adopt a vision grounded in concrete data, transparent costs and rigorous analysis of the real market. The clearest lesson is that the viability of green hydrogen depends less on a massive incentive policy and more on intelligent integration with the value chains where carbon is intrinsically difficult to eliminate — a perspective limited in scope but capable of generating measurable, sustainable impact. [1][2]

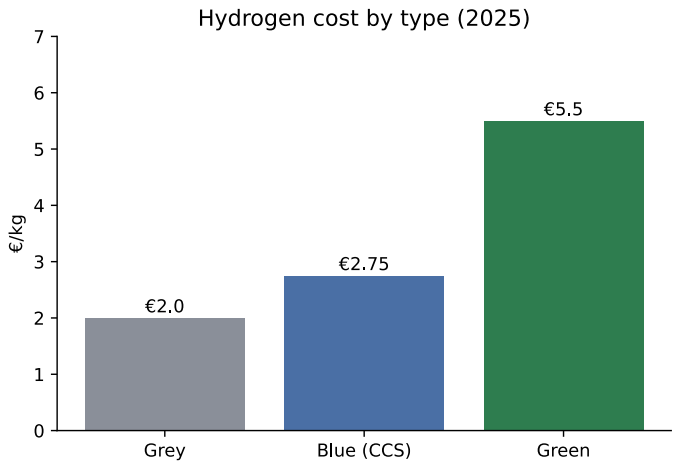


Figure 14.1 — Green hydrogen (€4-7/kg) is still far more expensive than gray or blue: the challenge is not technical, but one of cost and demand.

Where it does make sense

On the shore of Immingham harbor, where the Air Products towers stand outlined against the gray British sky, it is easy to imagine a group of engineers reviewing the flow diagrams of the green hydrogen that would feed a high-performance steel plant. The goal is not to celebrate a universal energy revolution but to validate a concrete bet: replacing natural gas and coal in processes where direct electrification is still unfeasible or prohibitively expensive. In this zone, green hydrogen becomes a precision tool, not a silver bullet promising a massive transition without foundations.

In the steel sector, direct reduced iron (DRI) using hydrogen offers a clear route to decarbonize a process that traditionally depends on coal. ArcelorMittal's initiatives to bring green hydrogen to its plants in Germany point in this direction, although the extra cost — up 25-40% since 2022, sitting between €4 and €7 per kilogram — remains a real barrier [2]. When integrated directly into the production flow, this premium can be offset by savings on emissions allowances, avoiding the need for large-scale electrical infrastructure to power a fully electric plant.

In fertilizers, ammonia remains one of the pillars of the world food system. Conventional production depends almost exclusively on natural gas as the source of gray hydrogen; replacing it with green hydrogen would drastically cut its CO_2 emissions. As already noted, the higher cost (between €4 and €7 per kilogram) can be offset in integrated plants by process-efficiency gains and emissions-reduction incentives, making production viable at regional scale. This application becomes one of the "no-alternative zones," where direct electrification cannot substitute for the intensive nature of the chemical process.

In fine chemicals and petrochemicals, green hydrogen is already used as a key reagent in hydrogenation processes that demand pressure and temperature conditions hard to reproduce with electricity. Integrating reactors fed by green hydrogen reduces the indirect emissions associated with transporting the gas and favors a cleaner supply chain. Locally produced green hydrogen reduces logistical risk and improves the traceability of the final product.

Finally, in heavy transport — long-haul trucks and cargo ships — green hydrogen and its synthetic derivatives offer a possible alternative where high-capacity batteries are not yet practical. This is, precisely, one of the uses the pipeline cut has postponed: projects like BP's in Duqm, canceled in December, have been left in suspense [1].

In short, green hydrogen will not die; it will simply find its place in the areas where substitution by electric alternatives is technologically impossible or economically unacceptable. The key is to treat it as a precision tool — a piece that fits perfectly within specific industrial processes — and not as a universal solution for every form of energy. This vision, grounded in real data and concrete cases, offers a clear path toward measurable, sustainable decarbonization.

Carbon capture and the bets still missing

At the dawn of 2026, carbon capture and storage (CCS) technology remains an incomplete promise. It advances little and remains expensive: the cost per tonne of CO_2 captured still far exceeds the economic expectations of most industries. This imbalance has turned CCS into a suspect refuge for actors seeking to prolong the useful life of fossil assets without facing the true challenge of reducing demand for hydrocarbon-based energy.

Carbon capture still accounts for a minimal fraction of global CO_2 emissions. The reality is that the technology works well in controlled scenarios — in refineries, for instance, or at large point sources of emissions — but runs into technical and regulatory barriers when applied to dispersed sources such as the cement or steel sectors. Current financing structures, moreover, penalize long-term projects: investors prefer quick returns, while the useful life of a secure underground reservoir stretches across decades.

This situation has fed a dangerous narrative: presenting CCS as the "silver bullet" for hard-to-abate emissions while using it as justification for delaying more forceful policies to move away from oil. The truth is that capture can be useful in specific niches — treating the emissions of steelmaking processes, or producing blue hydrogen with CCS — but only if combined with a real reduction in fossil-based energy and with clear policies limiting new carbon installations. It is no coincidence that BP, the same year it canceled Duqm, also paused its blue hydrogen CCS project in Whiting.

While CCS advances at a slow pace, other emerging bets are gaining prominence. Long-duration energy storage presents itself as a key solution for balancing intermittent renewables, with projects exploring energy conservation in salt caverns and other geological formations. Advanced geothermal, driven by new drilling and water-recirculation techniques, is beginning to make headway. And continuous electric grids — based on new-generation HVDC converters — help reduce transmission losses and ease the integration of renewable energy from remote areas into demand centers. None of these options is, on its own, a "silver bullet."

In this context, the approach must be holistic: CCS becomes a complementary tool for hard-to-abate emissions, not the core of a decarbonization strategy. Investment must prioritize technologies offering flexibility and scalability, such as long-duration storage and advanced geothermal, while regulatory policy is reinforced to prevent CCS from being used as a pretext to keep fossil infrastructure alive. Only then can we guarantee a genuine transition toward a sustainable, resilient energy future.

The promise and the peril

In industrial settings where decarbonization is a regulatory obligation and not merely a reputational goal, green hydrogen has become a key piece of the energy puzzle. In European steelworks, for example, direct iron reduction requires a heat source and a reducing agent that cannot easily be replaced by electricity without massive investment. Here, a steady flow of green H_2 — produced from wind or solar power — allows production to continue without increasing CO_2 emissions. Similarly, in fertilizer manufacturing, where ammonia concentrates a very large share of hydrogen consumption, the shift to green H_2 cuts the emissions associated with the process, as several pilot projects in the sector indicate.

This promise is reinforced, in some cases, by carbon capture and storage (CCS). Although the global rollout of CCS has advanced little, facilities integrated with blue hydrogen processes — where natural gas is reformed and the resulting CO_2 captured before release — can offer a pragmatic transition toward a cleaner sector. Combining green H_2 for the applications where direct electrification is not yet viable with CCS-fed blue H_2 for intermediate uses sketches an energy architecture that can cut emissions where no clean alternative exists today.

Yet the temptation to present this package as a universal solution threatens to lead us away from the cheapest, most efficient path. Most comparative studies show that directly electrifying processes such as road freight or residential heating is usually considerably cheaper than replacing them with green H_2, given current fuel costs (between €4 and €7/kg) and the associated infrastructure [2]. Moreover, the proliferation of "green" offers without rigorous certification has fed a market for hydrogen credits that behaves like the early carbon exchanges: the risk of greenwashing grows as companies seek to meet ESG targets with solutions that look simple but lack substance.

The danger, then, is not only economic but structural. Promoting green hydrogen and CCS as the sole answer can delay crucial investment in large-scale electrical infrastructure — ultrafast charging networks, battery energy storage — needed to make the most of the renewable surplus. Every year that capital goes

into building a large electrolysis plant rather than reinforcing the grid or developing demand-management systems is a lost opportunity to cut global emissions at a lower marginal cost.

Green hydrogen and CCS are indispensable pieces of the industrial puzzle: they offer viable solutions for sectors where direct electrification is not yet feasible. But their real value holds only if they are used in a targeted way, without pretensions of universality that could divert resources and delay the transition toward broader, more economically efficient electrification. The key is to balance the bets: prioritize electricity where it is most cost-effective, and reserve green hydrogen and CCS for the niches that truly need them.

The price of lost time

Picture a morning in Immingham, in the north of England: the towers of the Air Products plant, which a few years ago promised to be the lighthouse of green hydrogen, now stand silent behind a wall of cancellation notices. The project, along with some sixty similar initiatives, was withdrawn from the global program in 2025, cutting planned production by almost five million tonnes a year (−4.9 Mt/yr) [1]. This drop is not a mere numerical setback; it is a warning signal showing how excess hope can become a burden on the energy transition.

The collapse is explained by the mismatch between market expectations and technical and economic realities. The cost of green hydrogen has risen between 25% and 40% from its 2022 price, now sitting between €4 and €7/kg [2], far above gray hydrogen (€1.5-2.5/kg) and even blue hydrogen with carbon capture (€2-3.5/kg) [2]. In a context where investment grows more demanding and financing flows toward solutions with faster returns, green hydrogen remains a gray zone: promising, but too expensive for large-scale adoption without massive subsidies. This reality has also hit the big pipeline projects planned for 2030, which have fallen from 49 to 37 Mt — as we saw earlier — fragmenting the supply chain and narrowing opportunities in key sectors such as steel, fertilizers and chemicals.

But the true danger lies in the time lost. Every year that the adoption of green hydrogen and CCS infrastructure is delayed where it makes sense means more carbon emissions, higher decarbonization costs and a growing dependence on natural gas and fossil fuels. The lack of action amplifies the need for alternative solutions, such as direct electrification of industrial processes, which is proving ever more competitive both technically and economically.

In this scenario, deciding where and how to invest is crucial. We cannot allow the hype to divert us toward projects with long, uncertain returns; resources must be steered toward the niches where green hydrogen offers a clear advantage: decarbonizing processes that cannot easily be electrified, like iron reduction or ammonia synthesis. Simultaneously, research and development into more efficient, less costly CCS technologies must be driven forward, to achieve a realistic, balanced transition.

In short, if we fail to deploy these solutions with the urgency they demand, the cost of inaction will translate into higher energy prices, greater climate insecurity and a lost opportunity to position Europe's economies as leaders in clean technologies. This closes our review of the energy solutions. Part V carries this same urgency onto the terrain of trust: why detecting forgeries is not enough, and how the proof of existence is beginning to deploy on the first fronts — finance, photojournalism, healthcare, education and human rights.

---

[1] IEA, *Global Hydrogen Review 2025* / CleanTechnica. https://www.iea.org/reports/global-hydrogen-review-2025/executive-summary ↩

[2] IEA / The Extractor, *green H₂ cost 2022-2025*. https://cleantechnica.com/2025/07/07/hydrogens-brutal-month-billions-lost-as-mega-projects-collapse/ ↩

Detection Is Not Enough: The Proof of Existence as a Durable Defense

"For a successful technology, reality must take precedence over public relations, for nature cannot be fooled."

— Richard Feynman, Rogers Commission report on the Challenger, Appendix F (1986)

It was a cold January afternoon in 2026 when the publisher **LlumVerda** received, through its secure file-transfer channel, a video that promised to change the course of the local elections in Barcelona. The clip showed the leading candidate, Marc Soler, delivering an incendiary line: *"Taxes for culture will rise 30% this year."* The voice was clear, the facial expression looked genuine, and the backdrop of a municipal meeting room was instantly recognizable. The editorial team, aware of the exponential growth of deepfakes — from 500,000 in 2023 to almost **8 million** in 2025 according to Sumsub [9] — ran the file through *DeepSight Pro*, the most advanced manipulation detector available at the time. The program, powered by a neural network trained on the latest forgery databases, issued a terse report: **"Probably authentic,"** with a confidence score of 78%.

Wasting no time, the editor-in-chief, Laura Martínez, scheduled publication on the digital portal, adding a brief note reminding readers that visual truth could still be relied on. Traffic exploded: thousands of comments and shares appeared on social networks within hours, and the political debate was already catching fire.

Three days later, an internal alert from the IT security department revealed the truth. The video had been generated with a synthesis model more recent than any update to the DeepSight Pro detector: the technology for creating the fake had outrun the capacity to detect it. No provenance seal accompanied the file, and the detector — trained on forgeries of a previous generation — had been fooled. In other words, the reactive defense was always one step behind the attacker.

This scenario is no isolated fiction; the data from *DeepFake-Eval-2024* show that detectors failed in between **45% and 50%** of cases in real-world settings [1]. Whenever we depend on detection alone, we find ourselves trapped in an arms race where the generator advances faster than the analysis. The question that emerges is inevitable: **what if, instead of chasing the lie, we could prove the truth from its origin?**

The answer to that question defines the core of the durable defense proposed throughout this chapter. The point is not merely to detect a manipulation, but to **certify the provenance and integrity of content at the very moment it is created**, as the US authorities (NSA/FBI/CISA) have recommended since 2023 [4]. This vision, combining cryptographic proofs with blockchain anchoring — the file's fingerprint anchored on a public chain, with the RFC 3161 timestamp as a complement, as the *t2p* initiative applies it [7] — offers an answer to the danger of relying solely on detectors that grow more obsolete by the day.

In the sections ahead we will explore how this proof of existence translates into five sectoral fronts — compliance, photojournalism, healthcare, education and defense — and what concrete steps we can take today to turn theory into practice. The future of digital truth depends not just on seeing well, but on **knowing where things come from** and being able to prove it beyond doubt.

The race we are losing

The moment most organizations rely on a single line of defense — deepfake detection — they find themselves plunged into an asymmetric arms race that, so far, they have lost. The DeepFake-Eval-2024 benchmark, which gathers real samples from generator algorithms circulating on the market in 2024, reveals dramatic drops: the best open-source detectors record an accuracy of **45-50%** relative to laboratory settings, with performance of **0.75** on video, **0.86** on audio and **0.63** on images — all below the **~90%** a trained human analyst achieves [1]. Even the most advanced commercial tools reach only **78%** accuracy, a significant gap when the stakes are preserving trust in critical environments.

This gap is no accident. The phenomenon of *domain shift* — where models trained on clean, controlled data fail in real scenarios where audio and video are compressed, transcoded or manipulated by social platforms — creates a structural vulnerability that simple retraining cannot sustainably close. Every time a new compression filter is introduced or a recording algorithm changes, detection models lose precision

and must be recalibrated; meanwhile, deepfake generators advance at exponential speed, improving the quality and efficiency of their output. In this game of exchanges, the attacker decides *when* and *how* to strike, leaving reactive defenses at a constant disadvantage.

The volume of fake content only keeps growing. According to Sumsub data, the number of deepfakes in circulation went from **~500,000** in 2023 to **~8 million** in 2025 — a more than fifteenfold increase in just two years [9] — and this explosion amplifies the probability that a malicious piece of content reaches a wide audience before any detector can identify it. Moreover, the proliferation of accessible generation tools democratizes the power to create convincing forgeries, lowering the barriers to entry for malicious actors and multiplying the diversity of attack scenarios.

In short, depending exclusively on detection is like trying to secure a door with a chain while the burglar carries a saw. The technical limitations — domain shift, the lack of sustainable generalization — and the escalating volume and sophistication of deepfakes turn this strategy into a race that, so far, we have lost. To tip the balance, we must move toward solutions that do not merely identify the fake but **certify the origin and integrity of content** from the moment of its creation; that is the fundamental premise underpinning the five sectoral fronts we explore next.

From detection to provenance

In a Stanford University classroom, Professor Hany Farid shows a group of students an apparently authentic image of a political speech and, with a few clicks, reveals that the video contains a layer of manipulation invisible to trained eyes but not to current detection technology. "Reactive detection is like trying to put out a fire after the building has already collapsed," Farid says, stressing that **the key is knowing the origin of the media**, just as a nutrition label lists a product's ingredients [2]. The metaphor illustrates the limits of detectors — which fall to between 45% and 50% in real-world tests, according to the DeepFake-Eval-2024 study [1] — and, as we have seen, it points to a clear route: recording provenance from the very first instant.

The digital security expert Sam Gregory, founder of WITNESS, has championed for more than a decade the program **"Prepare, Don't Panic,"** a call to build an authenticity infrastructure that goes beyond mere detection [3]. In his presentations, Gregory describes how a chain of trust — signed capture, verifiable metadata and origin certification — can be integrated directly into multimedia production workflows. This infrastructure allows every file to carry with it immutable proof of the moment and context of its creation, preventing a deepfake from hiding its true history.

Even the most conservative government agencies recognize the need for this paradigm shift. The joint guidance from the NSA, the FBI and CISA — as noted, from 2023 — makes establishing provenance as high a priority as detecting manipulated content [4]. The document stresses that if AI generators advance faster than detection tools, betting everything on identifying the fake is a losing wager; **verifying where something comes from and whether it has been altered since origin** offers a mathematically verifiable, scalable defense.

This transition is not only technical but cultural. Moving from asking "Is it fake?" to asking "Where does it come from and how has it evolved?" means reconfiguring editorial, legal and business practice. Creators must adopt standards like **C2PA/Content Credentials**, despite the vulnerability of strippable metadata [5]; and organizations like OpenTimestamps are already sealing millions of files — around 750 million from the Internet Archive — with blockchain transactions, providing an irreversible anchoring layer [6].

In short, **provenance** becomes the core of a durable defense against deepfakes. It is not merely a complement to detection; it is the alternative that transforms the fight for truth into a verifiable, auditable process from the first instant of a piece of content's creation. This vision, backed by academic, industry and government experts, marks the turning point toward an era where trust rests on concrete proof rather than assumptions.

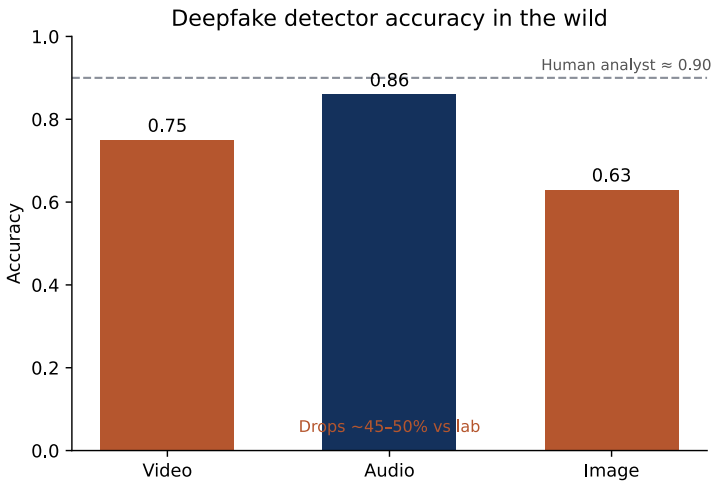


Figure 15.1 — Detector accuracy decays in the real world as generators improve; a provenance anchor does not degrade with time.

Two complementary pieces

In today's landscape, where deepfakes proliferate at exponential rates — from an estimated 500,000 cases in 2023 to nearly 8 million in 2025 alone, according to Sumsub [9] — detection by itself becomes an unsustainable arms race. The most advanced detectors fall to between 45% and 50% in real-world tests (DeepFake-Eval-2024) [1], as we have seen, while the generators of fake content keep outpacing technological improvements. A sustainable answer must go beyond "spot and block"; it must establish an unalterable chain of trust from the moment of creation, one that persists even if the metadata are deleted or tampered with.

C2PA / Content Credentials: the cryptographic record of origin

The C2PA consortium (Coalition for Content Provenance and Authenticity) has developed *Content Credentials*, a manifest signed with X.509 certificates that captures, as a JSON structure, the key information of the creative process: who generated the file, which tools were used and when each edit was applied. This infrastructure provides cryptographic proof of the "who/when" at origin, allowing editors, platforms and consumers to verify the provenance of an image or video without relying on a trusted third party [5]. By design, C2PA rests on digital signatures which, if kept intact, are impossible to forge.

The model, however, has a fundamental vulnerability: the credentials can be removed. Social networks, transcoding processes and even the simple export of a file often strip or overwrite this manifest. Once the credential disappears, no verifiable trace of the original content remains; its absence therefore proves nothing [5]. This limitation has led experts like Hany Farid to declare that "provenance is the best defense" — but only if it remains accessible and intact throughout the content's life cycle [2].

Proof of existence: the immutable anchor

The solution to this gap lies in the *proof of existence* mechanism described in chapter 13. The idea is simple but powerful: compute a cryptographic fingerprint of the file — a hash — and anchor it in a transaction on a public blockchain. This operation creates an immutable proof that the file existed at that moment, without needing to include any additional metadata; as a complement, a timestamp conforming to the RFC 3161 standard, independent of the chain, adds the certification of a qualified authority. Projects like OpenTimestamps have already applied this approach to millions of Internet Archive files with a single transaction, as we have seen (Peter Todd) [6].

Unlike C2PA, which depends on the preservation of the manifest, the hash anchor is independent of any intermediate layer. Even if the metadata are stripped or the file is reprocessed, the fingerprint anchored on the blockchain retains a verifiable anchor confirming the original content's unaltered existence on the recorded date [7].

Synergy between provenance and proof of existence

The two technologies are not rivals; they are complementary. C2PA contributes the "who/when" information at the moment of creation, while the fingerprint anchored on the public blockchain guarantees "this existed, unaltered, before this date." Combining them yields a robust defense that overcomes the limitations of each method on its own. Provenance can be lost, but the proof of existence persists; conversely, the proof of existence does not reveal who created the content, but the C2PA credential does — so long as it survives.

This double layer of verification transforms the fight against deepfakes from reactive detection into a preventive, auditable architecture. As the government authorities — NSA, FBI and CISA — recommend in their 2023 guidelines, establishing provenance is key to guaranteeing the integrity of information [4]. Without the immutable anchor of the proof of existence, that guarantee is left half-finished. Together, the two pieces form a durable, verifiable defense from the first instant of creation to the moment the content reaches the public's hands. This is the foundation on which the five sectoral fronts — compliance, photojournalism, healthcare, education and defense — can build policies and systems resilient to the age of deepfakes.

The unalterable "when": OpenTimestamps and RFC 3161

In the spring, while image-verification teams were being overwhelmed by the explosion of deepfakes — as we have seen [9] — a simple but powerful solution was gaining traction: verifiable proof of time. OpenTimestamps, the open standard devised by Peter Todd, turns any file into an immutable digital fingerprint by grouping thousands of hashes into a Merkle tree and anchoring its root to the Bitcoin blockchain with a single transaction. The procedure produces a compact proof that anyone can validate without relying on a centralized entity, as its design describes: "the proof is verifiable by anyone, without trusting any company" [6].

To date, OpenTimestamps has sealed a vast number of Internet Archive files with a single Bitcoin transaction, turning that immense historical catalog into an immutable record of its moment of existence [6]. Each Merkle tree root works as a cryptographic witness: if the content is modified, the resulting hash no longer matches the anchored root and the proof breaks. The mechanism thus offers a direct answer to the central problem of deepfakes: **when** a file first appeared, and the fact that it has not been altered since.

Yet, as Hany Farid warns, "provenance is the best defense" [2]; the timestamp, though a robust guarantee of the *when*, does not verify the authorship or the intrinsic truthfulness of the content. A video can be perfectly anchored to a specific date and still carry misleading or manipulated information. That is why the t2p model anchors the SHA-512 fingerprint, computed on the source device itself, directly in a transaction on a public blockchain — that anchor is the proof — and adds, as a reinforcement independent of the chain, an RFC 3161 timestamp (the eIDAS trusted-timestamping standard). This "belt and suspenders" ensures that even if the sealing authority's private keys expire or are compromised, the public proof remains verifiable thanks to the immutable fingerprint stored on-chain [7].

This architecture also allows smooth integration with emerging authenticity infrastructures like WITNESS, which promotes "Prepare, Don't Panic" and offers additional layers of verification based on signed metadata [3]. In practice, a content creator can generate the file's fingerprint at the moment of capture and anchor it on a public blockchain — directly, or aggregated into a Merkle tree through services like OpenTimestamps — adding, when the certification of an authority is needed, a complementary RFC 3161 seal. The result is a verifiable chain stating: *this file existed on this date, with no subsequent alterations*, which can be checked against the provenance metadata recorded by WITNESS or other regulatory bodies.

This combination of elements — local fingerprint, public-blockchain anchoring and a complementary RFC 3161 seal — constitutes the technological basis of the durable defense against deepfakes. The point is not merely to detect a manipulation, but to demonstrate unequivocally the original moment of creation and keep it intact across the digital life cycle. In an environment where AI generators evolve at accelerating speed, this proof of time becomes the cornerstone for the sectoral policies that will address the crisis of trust: from corporate compliance to photojournalism, through healthcare, education and defense. Without it, any detection would be an arms race; with it, we have a verifiable framework that turns **provenance** into a tangible, immutable asset.

The legal framework arrives

The entry into force of the **EU AI Act**, Article 50, on August 2, 2026 marks a clear turning point: any AI-generated content must carry a machine-readable label indicating its provenance and, specifically, deepfakes must be marked as "synthetic" [8]. In practice, companies already using **t2p**-style solutions — the fingerprint anchored on the public blockchain, with the RFC 3161 seal as a complement, as described in chapter 13 [7] — find themselves a step ahead: their infrastructure can automatically generate the provenance mark the regulation requires, turning the proof of existence into a direct answer to the legal obligation.

In the United States, the **TAKE IT DOWN Act**, signed in May 2025, raises the matter to the criminal level. The law explicitly criminalizes the non-consensual distribution of intimate images — including "digital forgeries," or deepfakes — and imposes on platforms an obligation to remove content within **48 hours** of a report [8]. The legal requirement is unambiguous: platforms must be able to demonstrate, with cryptographic evidence, that the content has been verified or removed in order to comply. This turns **provenance** into an operational requirement, not merely a good practice.

The final link in the legislative puzzle is the **DEFIANCE Act**, passed by the Senate in January 2026 and still awaiting a vote in the House of Representatives [8]. This law adds a civil avenue for the victims of digital manipulation, allowing them to bring claims based on the absence of verifiable evidence of a content's origin and integrity. In other words, a **proof of existence** — such as the one provided by OpenTimestamps, which has sealed millions of Internet Archive files [6] — becomes a decisive asset for sustaining any claim.

The common thread of these three initiatives is clear: the law no longer asks only "whether the content is fake," but **"where it comes from and how it can be proven to have been intact on a specific date."** Companies still relying solely on detectors — which, according to DeepFake-Eval-2024, fail in 45% to 50% of real-world settings [1] — will find themselves in a position of inevitable non-compliance. Those that have already integrated the proof of existence into their digital supply chains, by contrast, will be able to respond immediately with verifiable metadata and cryptographic seals, avoiding penalties and winning the market's trust.

So the legal framework does not merely push for the **labeling** of content; it establishes a **requirement of demonstration** that turns the proof of existence into an indispensable asset. Organizations that anticipate this need — like those adopting the C2PA standards, despite the limits of strippable metadata [5] — will be able to move toward a resilient trust architecture, avoiding the risk of being left trailing behind a regulation that advances at a steady pace. The answer is not just to detect; the answer is to **document**, **seal** and **prove** — the foundation on which the new era of digital credibility will be built.

Five fronts

The crisis of trust we have described is not an abstract catastrophe; it translates into economic losses, reputational damage, health risks and broken lives. Every sector shares a common logic: **detecting** a deepfake is only the first step, and often comes too late; the real defense consists of **documenting**, **sealing** and **proving** the provenance of content from the moment of its creation.

1. **Compliance and finance** — The most emblematic recent case is the video call that deceived Arup, causing a fraud of $25.6 million. Without unambiguous proof of the video's origin, the control teams were trapped in an arms race with the deepfake generators, and after-the-fact detection did not stop the improper payment. Here, the **proof of existence** would have made it possible to validate which audiovisual material was authentic before any transaction, blocking the scam in its infancy.
2. **Photojournalism and media** — The image remains the most powerful witness to the moment. When a manipulated photograph circulates as evidence of an event, the reputational damage to the paper or the journalist is irreversible. Sealing tools like OpenTimestamps have already **sealed around 750 million files** from the Internet Archive, offering a chain of trust that can be consulted to verify a photo's provenance before it reaches the public [6].
3. **Healthcare** — Manipulated medical images threaten diagnoses and treatments. An altered scan could lead to unnecessary surgery or, worse, to a tumor being missed. In this context, certified provenance is the best defense, as Hany Farid reminds us: "provenance is the best defense" [2]. An immutable record of every image — its fingerprint anchored on the public blockchain, with a

complementary timestamp — would guarantee that health professionals always work with original, verifiable data [7].

4. **Education** — Academic credentials and the authorship of digital works are constant targets of forgery. An altered digital diploma can open doors to unqualified employees, while a plagiarized or manipulated article can corrupt scientific knowledge. The proof of existence would allow every academic submission to be recorded with a secure timestamp, making retroactive revision of the content impossible without leaving a trace.

5. **Defense and human rights** — In armed conflicts, deepfakes become weapons of disinformation that can manipulate public opinion and falsify evidence of crimes. The ability to attribute with certainty the date and origin of a video or audio recording is fundamental for international courts and human rights organizations. Without this layer, the risk of "false narratives" becomes a direct threat to justice.

On all these fronts, the **logic** is identical: detection starts late, while proving from the origin arrives on time. In the chapters that follow we will explore how the proof of existence is integrated into each sector, revealing both the real threat and the concrete solutions for protecting ourselves in an era where the false can look indistinguishable from the true.

The promise and the peril

Imagine a global authenticity network where every video, audio file or digital document proudly carries an immutable seal certifying its creation date and the origin of the node that generated it. This layer of verifiable trust is as cheap as any low-fee blockchain transaction and can scale to millions of users without centralized intermediaries. The model proposed by **t2p** combines provenance — the information about who, where and when produced the content — with the cryptographic anchoring of the fingerprint on a public blockchain, reinforced with an RFC 3161 timestamp independent of the chain [7]. When an outlet comes under deepfake attack, the recipient can instantly consult the proof of existence and verify that the file already existed, intact, on a date prior to the alleged deception. In practice, this solution has been adopted by projects like OpenTimestamps, which has already sealed around **750 million** Internet Archive files with a single transaction [6], proving that the marginal cost is close to zero and the infrastructure can bear massive loads.

This promise is not merely theoretical; organizations like WITNESS, led by Sam Gregory, have already deployed an authenticity infrastructure that allows journalists and human rights organizations to "prepare, not panic" in the face of digital manipulation [3]. European regulation, with the AI Act (Art. 50) applying from August 2, 2026, explicitly requires AI-generated content to carry a machine-readable provenance mark to reduce the risks of deepfake dissemination [8]. This creates a legal framework that drives mass adoption of these technologies and gives the authorities a secure basis for acting on verifiable evidence.

The danger, however, lies in the potential **de-anonymization** of vulnerable creators. Hany Farid reminds us that "provenance is the best defense," but also warns that if too much information about the origin is revealed, investigative journalists, dissidents or victims of violence can be located and persecuted [2]. A badly designed system could turn provenance into a weapon against those who depend on anonymity for their personal safety; that is why provenance metadata, such as C2PA credentials, must be manageable with care [5].

Moreover, blind faith in the seal can breed an **excessive sense of veracity**: "if it carries a seal, it must be true" becomes a fallacy that neglects the other lines of defense. The proof of existence does not by itself detect fakes; it only demonstrates that the content existed at a given moment. Without active detection — and without media literacy that teaches people to question and cross-check — society could fall into a false security, believing the deepfake problem completely solved.

No technology, therefore, can be considered a silver bullet. The robust solution lies in the **combination** of reliable provenance, cryptographic anchoring, advanced detection systems and continuous public education about the limits and vulnerabilities of these tools. Only then can we harness the promise of verifiable trust without falling into the danger of unchecked control or of an overconfidence that, in the end, could prove as damaging as the deepfakes it seeks to fight.

Beth Noveck, in her book *Reboot*, underlines a reality that is often overlooked: deepfake detection systems are essentially an arms race in which the attacker can disable them with minimal crops or compressions of a few kilobytes, and the final result is always probabilistic, never certain [10]. This vulnerability makes exclusive confidence in detectors an illusion of security; as soon as a new generator algorithm appears, detection models are open to new forms of deception. The solution, therefore, is not to improve detector accuracy indefinitely, but to change the paradigm: to move from "detect whether it is fake" to "prove that this content existed, with a neutral anchor." Cryptographic anchoring on a blockchain — independent of the creator and beholden to no centralized entity — provides the immutable proof Noveck considers indispensable, closing the loop where detectors can only play the game of deception.

It's time to act

Digital platforms are the first line of defense: they must incorporate systems that preserve and display the verified provenance of every media file. That means integrating content credentials based on the C2PA model, but going beyond metadata that can be stripped; each file's fingerprint must be anchored to a public blockchain or to a reliable time registry like OpenTimestamps, which has already sealed around **750 million** Internet Archive files [6]. Every time a video or an image is uploaded, the platform must issue a "certificate of origin" and make it publicly accessible, allowing any user to trace the full path from the moment of creation to the last modification.

Creators and media must adopt a culture of publishing with proof of origin. It is not enough to publish a story; the chain of trust must be attached, demonstrating where the material was obtained, when it was created and which entity has digitally signed it. This approach follows the NSA/FBI/CISA recommendation (2023) to establish provenance as a key element in fighting manipulation [4]. Moreover, using the **t2p** framework — the fingerprint anchored on the public blockchain, with the RFC 3161 seal as a complement — media outlets can guarantee an immutable, verifiable proof without depending exclusively on vulnerable metadata [7].

Companies and institutions must anchor the fingerprint of their critical data before sharing it. Any financial report, clinical study or educational record must be signed and sealed on a decentralized network offering reliable timestamps. This not only protects against the malicious substitution of documents; it also satisfies the new regulatory obligations: the EU AI Act (Art. 50) takes effect on August 2, 2026 and demands traceability mechanisms for AI-generated content [8].

Regulators, for their part, must define interoperable standards that guarantee all stakeholders can verify provenance without technical barriers. That means adopting common specifications like C2PA, but also establishing audit and certification protocols for providers of timestamping services and public blockchains. International cooperation will be key to avoiding a "trust barrier" that stifles innovation while protecting citizens.

Finally, citizens must become the first filter: before believing or sharing a piece of content, its origin should be checked through accessible tools (browser extensions, mobile apps) that consult the credentials issued by platforms and media. This practice, however simple, can drastically reduce the spread of deepfakes, which have gone from 500,000 cases in 2023 to nearly 8 million in 2025 according to Sumsub [9].

In the face of the industrialized lie, proving the truth is the most powerful action; we can now turn the proof of existence into a durable defense. **Time2act** is not just a motto, but a call to every actor in the digital ecosystem so that, with provenance and timestamping, we can recover trust and confront deepfakes before they overwhelm us.

---

[1] arXiv — *DeepFake-Eval-2024: A Multi-Modal In-the-Wild Benchmark of Deepfakes.* https://arxiv.org/html/2503.02857v4 ↩

[2] ManageEngine Insights — *Content provenance is our best chance in the fight against deepfakes* (Hany Farid). https://insights.manageengine.com/artificial-intelligence/content-provenance-is-our-best-chance-in-the-fight-against-deepfakes/ ↩

[3] WITNESS — *Prepare, Don't Panic: Synthetic Media and Deepfakes* (Sam Gregory). https://www.witness.org/portfolio_page/sam-gregory/ ↩

[4] CISA / NSA / FBI — *Contextualizing Deepfake Threats to Organizations* (2023). https://www.cisa.gov/news-events/alerts/2023/09/12/nsa-fbi-and-cisa-release-cybersecurity-information-sheet-deepfake-threats ↩

[5] C2PA — *Content Credentials Explainer*. https://spec.c2pa.org/specifications/specifications/2.2/explainer/_attachments/Explainer.pdf ↩

[6] Peter Todd — *Carbon Dating the Internet Archive with OpenTimestamps*. https://petertodd.org/2017/carbon-dating-the-internet-archive-with-opentimestamps ↩

[7] ProofStamper — *File timestamp tools compared (RFC 3161 + blockchain anchoring)*. https://proofstamper.com/en/compare/file-timestamp-tools-2026/ ↩

[8] EU AI Act, Article 50 (transparency of AI-generated content). https://artificialintelligenceact.eu/article/50/ ↩

[9] Sumsub — *Global deepfake incidents surge tenfold*. https://sumsub.com/newsroom/sumsub-research-global-deepfake-incidents-surge-tenfold-from-2022-to-2023/ ↩

[10] Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 3. ↩

Money and Trust: Fraud, Compliance and Finance

"Sunlight is said to be the best of disinfectants; electric light the most efficient policeman."

— Louis D. Brandeis, Other People's Money (1914)

In early 2020, the manager of a Hong Kong branch of a Japanese company received a call with a voice he knew well: that of a director at headquarters. The director explained that the company was about to close an acquisition and that transfers worth **$35 million** had to be authorized. In his inbox were emails from the director and from a lawyer named Martin Zelner confirming the operation, and the manager began making the transfers. All of it was fake: the voice was a clone generated with "deep voice" technology. The United Arab Emirates investigation, revealed by a court document, identified at least **17 people involved**, with the funds scattered across accounts around the world, and traced **$400,000** to US accounts at Centennial Bank; the UAE ended up asking American investigators for help [8].

Four years later, the scam had learned to make video calls. In January 2024, a finance employee in the Hong Kong office of the British engineering firm **Arup** received instructions to transfer some **$25.6 million**; he suspected a phishing email, but a videoconference with the "chief financial officer" and several "colleagues" — every face and voice exact — reassured him, and he executed **15 transfers** to five accounts. Every participant on the call except him was a deepfake [1]. Chapter 87 dissects this case in depth; here our subject is the whole landscape: fraud, regulatory compliance and financial identity.

As Arup's CIO concluded after the incident: "*Copying a voice, an image or a video is within reach of anyone with little technical skill.*" If seeing and hearing are no longer enough to trust, finance needs another kind of proof.

When seeing is no longer believing

The digital age has turned trust into a matter of constant verification, and the deepfake phenomenon is accelerating that shift at a speed beyond any earlier prediction. The figures deliver an uncomfortable reality: Deloitte estimates that AI-generated fraud in the United States would go from $12.3 billion in 2023 to nearly **$40 billion** in 2027 [2]. This escalation is not just an economic projection; it is the synthesis of an ecosystem in which the voices and images of executives become easily manipulable.

In recent months, companies have reported **cloned CFO voices**, **synthetic videos simulating board meetings** and **"urgent" instructions sent outside official channels**. Each of these scams rests on human trust: facial recognition, vocal authenticity and familiarity with body language were considered unquestionably reliable until recently. Now all it takes is an image-generation algorithm or a voice-synthesis app to create a deception indistinguishable from the real thing.

The most corrosive impact lies not in the direct monetary losses but in the **systemic doubt** spreading through the entire financial environment. When an apparently legitimate call can be a deepfake, organizations must drastically expand their verification processes, generating longer waits to confirm payments, higher spending on authentication technologies and an administrative burden that dilutes operational efficiency. This friction does more than raise direct costs; it also erodes the speed of dealmaking, delaying investments and discouraging international collaboration.

Moreover, the cost of a failed verification becomes evident in cases like the Arup incident in Hong Kong in 2024, where on a video call every attendee — including the CFO — was a deepfake, leading to the transfer of **$25.6 million** without a reliable trace [1]; **as we have seen**, a single vulnerability can trigger massive losses and erode confidence across the entire banking system.

The authorities have begun to respond. The **EU AI Act**, with Article 50 requiring the labeling of deepfakes from **August 2, 2026** [5], and the **NIST SP 800-63-4** guidelines, which make injection-attack detection mandatory at the higher assurance levels for biometric processes [4], are essential steps toward a regulatory defense. Yet these measures only mitigate the problem if institutions adopt a trust architecture that goes beyond simple visual or auditory recognition.

In sum, the proliferation of deepfakes turns every financial interaction into a scenario of uncertainty. The hidden cost — the time and resources devoted to validating the truthfulness of each communication —

becomes a structural burden threatening the fluidity of the global market. To preserve the stability of the financial system, we need a response that re-establishes trust not in appearances, but in the verifiable, immutable origin of every transaction.

The KYC front: injection attacks

At a bank registration office in New York, an operator received an application from a new customer that at first glance looked entirely legitimate: a scanned identity document and a recent selfie. The verification flow broke down when an invisible layer of software — a virtual camera injected in real time — replaced the device's image with an AI-generated version, producing a "selfie" perfectly aligned with the document but which was, in reality, an advanced deepfake capable of beating the most sophisticated biometric algorithms.

iProov records a drastic rise in injection attacks and biometric bypasses against KYC processes. One report describes how a deepfake opened dozens of fraudulent accounts at a bank, exploiting the vulnerability of the "document + selfie" model [4]. The attack does not merely fool facial recognition; it manipulates the entire data-input chain, so that every step — image capture, compression, transmission and analysis — takes place on a falsified foundation. The emerging technical response starts from a fundamental premise: integrity must be bound to the moment and the point of capture, not just to the final file. The NIST SP 800-63-4 standard already mandates, at the highest assurance level, the detection of injection attacks [4]; that is why systems must establish a "signed capture," a cryptographic record inserted at the instant the physical sensor captures the image, so that any subsequent manipulation — including the addition of a virtual layer — becomes evident because the hash of the original capture has already been archived on an immutable chain.

This logic underpins the proof-of-existence infrastructure (t2p) of *time2act*. Instead of relying solely on the document's content, the system generates a temporal, verifiable fingerprint stored in a distributed registry; every selfie, scan or biometric signal becomes an immutable transaction, so that any injection attempt would have to rewrite the blockchain and would be detected immediately.

The impact is not limited to reducing fraud: it also answers recent regulatory alerts, such as FinCEN's notice of November 13, 2024 on the need to strengthen CDD controls against deepfakes [3] and the EU AI Act Article 50 obligation to label this type of content from August 2, 2026 [5]. With secure capture and a verifiable registry, financial institutions can meet these demands without bolting on any onerous operational layer.

In sum, identity verification has moved from paper and screen to the moment the signal is generated. The combination of mandatory injection detection (NIST) and immutable proofs of existence (t2p) creates a barrier that, so far, fraudsters have not been able to cross; the challenge is to deploy this architecture at global scale, with trust rooted in verifiable origin, not in deceptive appearances generated by AI.

The regulator moves

As we have seen, on November 13, 2024 FinCEN issued an alert that rang like a thunderclap through boardrooms and compliance departments around the world: deepfakes and generative AI were becoming the preferred tools for evading identity verification and customer due diligence (CDD) controls [3]. The notice did more than describe a new threat; it demanded an immediate change in how institutions record and report suspicions. FinCEN asked that related suspicious activity reports incorporate the key term "FIN-2024-DEEPFAKEFRAUD," creating a uniform marker for subsequent analysis and cross-border cooperation. This request, however simple in appearance, obliges financial entities to overhaul their detection systems, integrate advanced authentication technologies and establish audit processes capable of tracing the real origin of an image or an audio recording.

As already noted, Article 50 of the EU AI Act establishes that from August 2, 2026 all AI-generated content — deepfakes included — must be clearly labeled as such [5]. The label is not merely a visual warning; it entails the insertion of verifiable metadata readable by automated compliance systems, facilitating early detection and the choice of corrective action. This obligation creates a legal framework that turns transparency into a regulatory requirement, not a business option.

Furthermore, the new eIDAS 2.0 has introduced two key components: Qualified Electronic Ledgers (QEL) and Qualified Electronic Archiving (QEA). These certified electronic registries offer an immutable audit

trail that can be validated at the regulatory level [6]; every transaction, KYC document or suspicion report is recorded in a qualified ledger, and its file is preserved in an electronic archiving environment that guarantees integrity and inalterability over time. This infrastructure allows the authorities to verify not only that a piece of information was submitted, but also that it comes from an authentic source and has not been manipulated afterward.

The message that emerges is clear: compliance can no longer be limited to "saving PDFs." The mere preservation of scanned or conventionally digital documents offers no guarantee against the sophisticated simulations created by AI. The objective shifts, as we have seen with NIST SP 800-63-4, to demonstrating — verifiably and auditably — the integrity and origin of a piece of information from the moment of its creation, through mandatory injection detection, immutable proofs of existence like t2p and the QEL/QEA regulated by eIDAS 2.0, forming a barrier that is nearly impassable for fraudsters.

This regulatory evolution does more than raise the level of security; it creates an ecosystem in which trust rests on verifiable origin rather than on deceptive appearances generated by algorithms. Companies that adopt these advanced practices will be able to transform their compliance model from an administrative chore into a proactive defense against emerging threats, becoming pioneers of a new paradigm of digital trust.

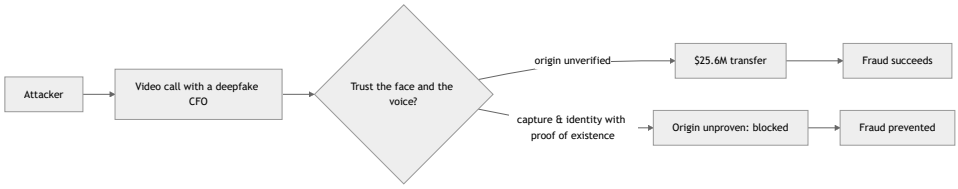


Figure 16.1 — In deepfake fraud, trust breaks at identity and instruction; a proof of existence anchored at the origin adds an independent verification point.

Evidence that can be fabricated

In the contemporary judicial world, the chain of custody is the pillar on which digital evidence rests: metadata recording every handling, witnesses certifying the origin and documents guaranteeing integrity. This model, born in the era of paper and electronic formats, is now suffering deep erosion at the hands of deepfakes; when a video or an audio recording is generated entirely by AI, merely presenting metadata is no longer enough, and confidence in authenticity collapses before an attack capable of reproducing with precision any known voice or face.

In Hong Kong, as we have seen, a 2024 fraud operation replaced a "CFO" and several executives with AI-generated avatars on a video call that authorized the transfer of a large sum to external accounts. The transaction was validated because the participants believed they were facing real identities, when all of them were deepfakes [1]. This case illustrates how forgery strikes not only finance but the judicial process: if a piece of audiovisual evidence can deceive, any trial built on it becomes vulnerable.

In the United States, the legislative response is advancing toward a more rigorous framework. The **Federal Rules of Evidence 707 (FRE 707)** — newly introduced — set specific criteria for machine-generated evidence, requiring proof of its reliability and of the impossibility of subsequent manipulation [7]. In parallel, the proposed rule **901(c)** seeks to establish a standard for potentially fabricated or altered electronic evidence, obliging lawyers to present verifiable proof of origin before a judge can accept it as valid [7].

These new provisions turn the proof of existence — a cryptographic hash anchored with a timestamp — into an essential requirement for any document or image to be presented in court. A hash, generated at the moment of creation and recorded on a public blockchain, provides an immutable digital fingerprint: no later alteration can change it without being detected. E-discovery audits thus come to depend not only on collecting files, but on certifying that those files existed at a given instant and that their content has not varied since.

The impact is reinforced by the emerging regulations; the EU AI Act requires any AI-generated content to be labeled as a "deepfake" from August 2, 2026 [5], and the eIDAS 2.0 framework introduces Qualified Electronic Ledgers and qualified archiving, which guarantee tamper-proof auditing for critical electronic documents [6].

As we have seen, the combination of stricter rules (FRE 707, rule 901(c)) and the NIST SP 800-63-4 technical standard, which requires injection-attack detection at the higher levels [4], together with blockchain-based proof-of-existence solutions, creates a new compliance infrastructure. It protects not only finance but the integrity of the judicial system, turning verification from a reactive task into a proactive defense against the fabrication of evidence; the key is that every document, image or audio file can be traced to its origin and corroborated with an unalterable timestamp.

The proof of existence as compliance infrastructure

When an organization decides that trust must begin **at the origin**, the first step is to capture every identity datum with an immutable digital signature verifiable in real time. During onboarding, the KYC is captured from the user's device through a public-key certificate that generates a hash of the presented document (passport, ID card or biometric selfie). This hash, called the *proof of capture*, is signed on the spot with a private key held in the Secure Enclave and bound to a timestamp conforming to RFC 3161; the result is published on a public blockchain that guarantees **resistance to the injection attack** described by iProov and required by NIST SP 800-63-4 [4]. Any attempt to substitute the biometric image with a virtual camera is thus detected: the hash stored on the chain matches no other version of the file.

Once the capture is registered, an **immutable audit trail** is generated combining the original hash, the timestamp and the blockchain anchor. This record is accessible to auditors and regulators without exposing the sensitive content, since only the cryptographic proofs (hashes and signatures) are shared. When an entity must demonstrate the truthfulness of a document in court, it can present this set of metadata as **evidence of provenance**. American courts, for example, already have rules like FRE 901/902 requiring a reliable chain of custody; the *proof of existence* offers exactly that: a verifiable link between the original document and the immutable record, defeating the "it's a deepfake" argument that has fed cases like the **$25.6 million** fraud in Hong Kong, where every attendee was an AI-generated avatar [1].

The proof of existence also guarantees **integrity and a certain date** for regulatory documents and contracts. Each version of the document is registered with its hash and timestamp; any later modification generates a new record that leaves clear evidence of the change. This aligns perfectly with the requirements of the Qualified Electronic Ledgers and Qualified Electronic Archiving introduced by eIDAS 2.0, which demand tamper-proof auditing for high-value electronic transactions [6]. The solution therefore not only complies with European regulation — it **amplifies** it, offering a decentralized verification layer that national authorities can consult without depending on a single provider.

It must be stressed that this mechanism **does not replace** traditional controls like the "four eyes" principle or out-of-band communications; rather, it acts as a **complementary layer** reinforcing every step of the process. The file never leaves the device: the sensitive information stays local, protected by the operating system's encryption, while only the cryptographic proofs travel to the network. **Privacy** is thus maintained without sacrificing verifiability.

According to Deloitte, as we have seen, the figures are alarming [2]; the proof of existence turns trust into tangible infrastructure: capture signed and sealed at the origin, immutable trails that support litigation and audits, integrity certified under eIDAS 2.0 and a privacy-first architecture. The decision is clear: integrate *t2p* right now, or face a future in which trust dissolves with every new digital manipulation.

The promise and the peril

Imagine a meeting room in Hong Kong where, with all the fluency of a normal video call, transfers totaling $25.6 million are signed off. The "CFO" presenting the documents is a deepfake so convincing that no one suspects its digital origin. The operation, as described by Arup, shows how trust in a visual identity can be exploited and turned into executive fraud at global scale [1].

This scene displays the central promise of *t2p*: a layer of integrity that ties every financial instruction to its verifiable origin, regardless of visual or audio format. With proofs of existence recorded in qualified ledgers — under eIDAS 2.0 — a single cryptographic query can demonstrate that the document has not been altered since the moment of its creation and that the issuer is who they claim to be [6]. The result is more reliable onboarding, where entities can validate KYC in minutes instead of hours or days, drastically cutting operating costs and eliminating the margin of human error. Audits become processes of instant

verification: an immutable chain of custody that can be presented in court as admissible evidence, overcoming the limits of the traditional FRE 707 rule and the proposed 901(c) [7].

The danger, however, lies in the illusion that a technological layer can entirely replace internal controls and human judgment. NIST SP 800-63-4 advocates mandatory injection-attack detection at the higher levels, recognizing that virtual cameras can deceive biometric systems if they are not properly protected [4]. Attackers are already experimenting with false-data "injections" and biometric bypasses to evade KYC verification, as the iProov report on emerging vulnerabilities shows [4]. As we have seen, the exponential rise of AI-generated fraud — according to Deloitte, from $12.3 billion in 2023 to nearly $40 billion in 2027 — is not solved by a proof of existence alone; the technology creates new attack vectors that demand constant oversight and robust governance policies [2].

A badly designed implementation, moreover, can compromise customers' privacy. Storing detailed metadata in qualified ledgers implies potential exposure if the encryption keys are poorly managed or if sensitive information is shared without the consent mechanisms the EU AI Act requires; as we have seen, that law mandates labeling and disclosure of deepfakes in commercial contexts [5]. The solution is therefore not only technological but legal and ethical: clear protocols for handling personal data must be established, and proofs of existence must be "zero-knowledge" where necessary.

The proof of existence represents an essential layer within a broader architecture of trust. It is not a silver bullet that eliminates every risk, but a mechanism that drastically reduces fraud and restores the credibility of financial commerce. Its effectiveness depends on how it is combined with rigorous internal controls, regulatory oversight and responsible privacy practices; only then can we balance the promise of a safer financial ecosystem against the danger of excessive confidence in a single technological solution.

It's time to act

Banks and fintechs are the first shield against the deepfake storm that has already shown its destructive potential; the Arup case, in which some **$25.6 million** was transferred through a video call populated exclusively by AI-generated avatars, is a clear warning that trust in identities and instructions can no longer be taken as safe [1] — as we have seen. To cut off this vulnerability, the **KYC must be anchored**, along with all its audit trails, within an out-of-band verifiable framework, especially for payment orders above a few thousand dollars. That means integrating solutions based on **Qualified Electronic Ledgers** and **Qualified Electronic Archiving**, as provided by eIDAS 2.0, which guarantee an immutable chain where every step of the customer-onboarding process is recorded with verifiable cryptographic seals [6]; even if a malicious actor tries to fool the system with a forged biometric image, the decentralized registry will expose the mismatch between the declared origin and the stored data.

Companies handling contracts and sensitive documents must **digitally seal these assets before sharing them**, using blockchain-based proofs of existence (t2p). A cryptographic seal not only certifies the document's state at the moment of signing; it also provides verifiable, immutable proof that any interested party can consult without blindly trusting the original issuer. This practice becomes especially crucial as the courts begin to address potentially fabricated evidence through rule 901(c) and FRE 707, where the traditional chain of custody is already in question [7].

Auditors and lawyers must become **champions of proofs of existence** in their professional practice. It is not enough to check that a document exists; its integrity must be tied to an independent verifiable proof, stored in a qualified electronic registry. That means auditing the content as well as the creation metadata and digital signatures, ensuring that any injection attempt or biometric bypass — as NIST SP 800-63-4 describes — is detected in real time [4].

Regulators must lead the definition of **interoperable standards and legal recognition of the cryptographic seal**, establishing frameworks that oblige financial institutions to adopt blockchain-based solutions as an integral part of their compliance programs. The application of the EU AI Act (Art. 50) on August 2, **as we have seen**, is a crucial step toward regulating these emerging risks [5]; but we must go further, legally recognizing proofs of existence as admissible evidence and establishing **out-of-band verification** requirements for large transactions.

In short, trust is the financial system's foundational asset; proving it is not a luxury but an urgent necessity. Every bank, fintech, company, auditor, lawyer and regulator has a concrete, complementary role in this ecosystem: capture, seal, verify and recognize proofs of existence as the cornerstone of credibility.

Proving trust is time2act — the moment to put these measures into practice is now, before next-generation deepfakes become the norm rather than the exception.

---

[1] CNN Business — *Finance worker pays out $25 million after video call with deepfake 'CFO'* (Arup case). https://www.cnn.com/2024/05/16/tech/arup-deepfake-scam-loss-hong-kong-intl-hnk · World Economic Forum. https://www.weforum.org/stories/2025/02/deepfake-ai-cybercrime-arup/ ↩

[2] Deloitte — *Generative AI is expected to magnify the risk of deepfakes and other fraud in banking* (≈$40 billion in 2027). https://www.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html ↩

[3] FinCEN — *Alert on Fraud Schemes Involving Deepfake Media* (Nov. 13, 2024). https://www.fincen.gov/system/files/shared/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf ↩

[4] iProov / Biometric Update — *Biometric injection attack surge* (KYC, NIST SP 800-63-4). https://www.biometricupdate.com/202604/biometric-injection-attack-surge-spreads-to-ios-iproov-report ↩

[5] EU AI Act, Article 50 (deepfake labeling; applies Aug. 2, 2026). https://artificialintelligenceact.eu/article/50/ ↩

[6] Signaturit — *eIDAS 2.0: Qualified Electronic Ledgers and Archiving*. https://www.signaturit.com/blog/eidas-2-regulation/ ↩

[7] Thomson Reuters Institute — *Deepfakes on trial: evidence authentication (FRE 707, 901(c))*. https://www.thomsonreuters.com/en-us/posts/ai-in-courts/deepfakes-evidence-authentication/ ↩

[8] Forbes (Thomas Brewster) — *Fraudsters Cloned Company Director's Voice In $35 Million Heist, Police Find* (14-10-2021; UAE court documents: at least 17 people involved, $400,000 traced to Centennial Bank accounts in the US). https://www.forbes.com/sites/thomasbrewster/2021/10/14/huge-bank-fraud-uses-deep-fake-voice-tech-to-steal-millions/ ↩

The Image as Witness: Photojournalism and Media

"A photograph passes for incontrovertible proof that a given thing happened."

— Susan Sontag, On Photography (1977)

It was the morning of May 22, 2023 when a digital whistle tore through social networks like an invisible explosion. An account marked as "verified," posing as a well-known news outlet, shared a photograph: a column of fire and smoke rising violently in the shadow of the Pentagon. The image spread within seconds; the algorithms pushed it into the timelines of thousands of users and, for a few minutes, it even briefly moved the North American financial markets before anyone could confirm anything.

The comments multiplied like fireworks: "Bombing at the Pentagon!", "Terrorist attack!" and, among the reactions, the key question that still echoes today: **where does this image come from?** The answer, at that instant, was a void. No one could trace the file's origin. The "verified" account was deleted shortly after the first wave of alerts, leaving only the viral image and a series of rumors running faster than the truth.

Amid the chaos, Hany Farid, a pioneer of digital forensic analysis, examined the photograph and concluded that, in all likelihood, it had been generated by artificial intelligence. The image showed the inconsistencies typical of generative models and, above all, could not be tied to any source or any known camera [1].

While financial analysts recalculated and market regulators weighed the impact of visual disinformation, the world of journalism found itself facing a scenario that would have seemed impossible only two decades earlier: a single image, without any verifiable trace, could trigger stock market movements and change public perception in a matter of minutes. The uncertainty about that photograph's origin was not merely a technical problem; it was a direct threat to the credibility of journalism and to democratic debate.

That brief digital explosion thus became the first testimony of a new era: one in which **verifiability** — the power to prove where each image comes from — becomes the only shield capable of protecting society from the chaos a single pixel can breed. The need to sign content at the origin, and to anchor it in an immutable proof of existence, was no longer just a technical proposal; it was an urgent condition for the survival of public discourse.

The image is no longer proof

For more than a century, photography enjoyed an aura of incontestable truth: the black and white of a war front, the smile of a family that becomes a historic icon or the fire of an explosion capturing the decisive instant of a conflict. That "frozen moment" was considered reliable testimony because the technology of the age demanded a physical, visible chain — film, laboratory, print — that made it almost impossible to manipulate without leaving a trace. Trust in the image was built on the presumption that its origin was evident and immutable.

Today that premise has come apart. Generative neural networks, fed by models trained on billions of images, can create frames indistinguishable from originals in a matter of seconds. As we saw in May 2023, the fake image of the "Pentagon explosion," an AI-generated composition spread by an account proclaiming itself a news outlet, caused a brief but significant swing in the financial markets. Hany Farid's analysis showed how digital manipulation can deceive even the most experienced observers [1].

The danger, however, does not lie only in "pure" deepfakes. In the 2024 elections, the media observed that "cheap fakes" — real videos or photos pulled out of their original context or retouched with basic tools — outnumbered AI deepfakes by a ratio of roughly seven to one [2]. These manipulations, though not fully synthetic, were just as devastating, because their credibility depended on knowing their origin and moment of capture. A clip of a political speech re-edited to add a compromising sentence, or a photograph of a demonstration shifted in time, can generate the same impact as a perfectly recreated digital avatar.

This reveals the fundamental truth: the central problem is not so much creating fake images as **provenance** and **context**. Who took a photo, at exactly what moment, with what technical settings and

whether it has been altered afterward constitute the only criterion for judging its truthfulness; without that information, any pixel becomes a potentially deceptive datum.

The solutions from the field of verifiability — content credentials (C2PA) and immutable proofs of existence — aim precisely at this gap. The digital signature at the moment of capture, already implemented in cameras like the Leica M11-P or the Nikon prototype that validated C2PA certification in February 2025 [6][4], guarantees a chain of authenticity from the first instant. Yet the metadata attached to these credentials can be stripped or altered by social platforms and transcoding processes, leaving the image without a verifiable trace [7].

Anchoring in a blockchain-based proof of existence provides an additional layer of security: an immutable hash recorded on the chain ensures that, even if the metadata disappear, the digital testimony persists and can be verified by any actor. This double defense — signature at the origin and immutable anchor — constitutes the only viable shield in an era when the pixel becomes an instrument of confusion as much as of information.

Content credentials: signing at the origin

At a moment when visual truthfulness has become a scarce good, the industry's answer has been to build an infrastructure that restores trust from the first click of the shutter. The C2PA project (Coalition for Content Provenance and Authenticity) and its Content Credentials initiative offer a cryptographic framework for recording the origin and the entire editing history of an image, so that every step — from the initial capture to later modifications — is documented in a signed manifest. This proposal does more than provide transparency; it creates a verifiable audit trail for any actor, from the journalist to the final reader.

As we have seen, the Content Authenticity Initiative (CAI) passed 5,000 members in 2025, including major digital platforms, news agencies and hardware manufacturers [3]. This critical mass reflects a shared awareness: proof of authenticity can no longer be a luxury reserved for the best-resourced organizations, but a basic infrastructure for the entire flow of information.

The qualitative leap into hardware has been decisive. The **Leica M11-P** launched as the first camera to incorporate Content Credentials directly in a secure chip, capable of signing each photograph at the exact moment of capture. That cryptographic signature is embedded in the file's metadata, ensuring that any later attempt at manipulation is detected. The innovation has not stayed in the premium range; Nikon, Canon, Sony, Samsung and even Google Pixel have followed suit, integrating similar chips into their most recent lines [6]. The result is an ecosystem where visual trust becomes a product feature as much as resolution or shutter speed.

At the editorial level, the impact has been remarkable. AFP and Nikon validated in February 2025 a prototype camera that **signs photos with C2PA at the moment of shooting**, guaranteeing a chain of authenticity from the source [4]. As we have seen, France Télévisions adopted the C2PA signature for its daily news programs at the end of 2025, and the IBC Accelerator "Stamping Your Content" — which brings together the BBC, the EBU and the Associated Press — has created open tools so that newsrooms can sign and verify content credentials without technical friction [5].

This convergence of hardware, open protocols and editorial workflows constitutes the integral response the industry needs to counter the age of "cheap fakes." It is not just an additional security layer; it transforms the image's life cycle, with proof of authenticity at every step from capture to distribution. Yet, as the C2PA framework itself stresses, metadata can be stripped or altered during recompression and sharing on social networks; that is why anchoring this information on a public blockchain becomes indispensable to guarantee that the proof of existence survives any attempt at suppression [7].

In short, the combination of C2PA credentials with immutable proofs of existence represents the road to a new visual trust. The industry has shown that it is possible to create an ecosystem where truthfulness is not an option but a guarantee built into the product and the editorial flow — key to preserving a public debate grounded in verifiable facts and to confronting the emerging threats of generative AI.

The hole: metadata get erased

When a photograph leaves its original environment — the camera, the newsroom lab or the app that captured it — it enters a digital ecosystem dominated by social networks and storage services. On this journey, the metadata carrying the C2PA signature often meet an invisible barrier: the upload and transcoding process. Platforms like Instagram, TikTok or WhatsApp reprocess every image to optimize size and bandwidth; that reprocessing strips the layers of information not essential for display — including the authenticity credentials.

This practice is more than a technical inconvenience; it is a structural vulnerability. An image shared in a group chat or posted in a social media thread can lose its Content Credential entirely, without any notice to the creator or the viewer. The disappearance of the signature does not necessarily mean the photo is fake — a photograph without a credential can be entirely genuine — but it also prevents any subsequent verification from confirming its original who and when. In other words, the absence of a metadata layer leaves an interpretive void: the reader faces a visual testimony stripped of its chain of custody.

Academic critiques have underlined this point, noting that the C2PA specification, on its own, does not offer sufficient robustness against the loss of metadata in the contemporary digital flow [7]. That means the "who and when" — the fundamental elements of visual testimony — can vanish at exactly the moment they are needed most: during an information crisis, a journalistic investigation or a public debate.

This hole is not merely technical; it has social and political repercussions. In the era of "cheap fakes," where badly contextualized videos outweighed AI deepfakes by about seven times in impact [2], the capacity to prove that an image comes from a reliable source becomes a critical asset. If the credentials disappear, the debate can drift into suspicion and skepticism, feeding the polarization this book seeks to counter.

That is why the solution is not to reject C2PA, but to recognize its limit and complement it with an immutable proof-of-existence layer. A blockchain-based anchor can record the hash of the image and its credentials at the moment of capture, ensuring that even if the metadata are erased during distribution, the permanent record persists and is verifiable by any actor in the ecosystem. This double protection — the C2PA signature for immediate authenticity and the proof of existence for long-term persistence — constitutes a robust defense against the manipulations of generative AI and guarantees that visual testimony **remains** reliable, even when the digital road turns uncertain.

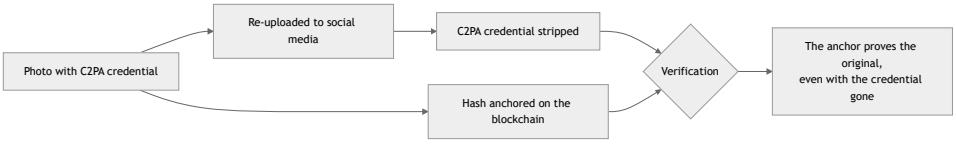


Figure 17.1 — When an image is re-uploaded, the C2PA metadata can disappear; the hash anchored to the public chain still proves the original existed, unaltered.

The anchor that survives

In a world where visual trust has become a scarce commodity, the **t2p** model offers the answer journalism needs to preserve the truth of an image from the moment it is captured. The key is not only signing it with C2PA credentials — an immediate authenticity layer that has already mobilized more than 5,000 global members and been adopted by entities like France Télévisions and the IBC Accelerator — but anchoring its hash to an immutable public chain. This double protection guarantees that even if the metadata are erased or altered in the distribution phase (as the studies on the vulnerability of C2PA metadata warn [7]), the fundamental element of the proof — the hash — persists and can be verified by any actor in the ecosystem without intermediaries.

At AFP, the editorial workflow incorporates an encrypted watermark that directly links the published photo to the original file stored in its secure infrastructure [4]. This practice illustrates how the combination of **C2PA signature** and **proof-of-existence anchor** turns a simple capture into verifiable testimony from the first instant. When the journalist, in the field or in the studio, shoots with a Leica M11-P — the first camera to carry Content Credentials in its secure chip — the image leaves the viewfinder already signed and, simultaneously, its hash can be sent to a public blockchain where it receives an immutable timestamp. At that moment a "publish with proof" version is created, consultable by editors, social platforms or auditors at any time.

This architecture solves the double problem exposed by cheap fakes during the elections — where badly contextualized videos outnumbered AI deepfakes by a ratio of roughly 7:1 [2] — by guaranteeing that provenance and temporal integrity are unalterable. Against incidents like the fake AI-generated image of the "Pentagon explosion" in May — which briefly shook the markets until Hany Farid dismantled it [1] — journalism now has a tool capable of demonstrating that an image is not only authentic at the moment of capture, but has been immutably recorded and can be traced to its origin.

The joint recommendations of the NSA and CISA in early 2025 on the use of Content Credentials [8] reinforce the idea that visual security is a matter of national interest, where the proof of existence acts as a universal code of trust. Any outlet — from major television networks to independent blogs — can adopt this framework without massive investment: it only takes integrating the t2p client, which generates and registers the hashes.

Thus the combination of C2PA and t2p creates a double barrier — an immediate authenticity layer and a long-term immutable anchor — that protects images against the erosion of trust caused by generative AI, restoring photojournalism's role as reliable testimony and giving journalists a practical tool for publishing with proof.

The institutions fall into line

When, in early 2025, the National Security Agency (NSA) and the Cybersecurity and Infrastructure Security Agency (CISA) published their guidelines recommending the use of Content Credentials — the operational framework of the C2PA initiative — a genuine convergence between national security and journalism took place. In a joint document, "Strengthening Multimedia Integrity in the Generative AI Era," the agencies warned that media integrity is a critical defense against AI-based manipulation and recommended adopting cryptographic signing at the moment images and videos are captured or published [8]. This recommendation, though initially aimed at federal entities, has produced a cascade effect: the big news networks, streaming portals and even social networks have begun integrating the C2PA protocol as a minimum requirement for visual content.

At the same scale, the JPEG Trust consortium, responsible for the ISO/IEC 21617-1 standard, has widened its scope by extending the C2PA engine inside the JPEG specification. This means that any JPEG file generated by a compatible camera can contain, by default, a verifiable digital signature and a record of authenticated metadata [8]. The move from an optional add-on to an integral component of the world's most widely used format consolidates the credibility of Content Credentials and elevates their status to an international standard. Now the simple act of creating an image with a modern camera like the Leica M11-P, or the Nikon, Canon and Sony models that already integrate secure chips, automatically produces a C2PA certificate embedded in the file [6].

This regulatory evolution coincides with the entry into force of the EU AI Act, which imposes a labeling obligation for all content generated or modified by AI. Article 50 of the regulation requires platforms and providers to mark clearly whether visual material has been created or manipulated by algorithms. In this context, the proof of C2PA credentials becomes a direct answer to the legal obligation: the metadata signed at the moment of capture acts as a "trust label" and helps meet the AI Act's transparency requirement. What was a good practice — using C2PA to raise credibility — thus turns into a regulatory expectation.

The pragmatic effects of this alignment are immediate. Media companies that have already adopted the credentials, as we have seen with France Télévisions in 2025 [5] or AFP with its C2PA-certified Nikon prototype [4], are betting on reducing incidents of forged-image distribution. Moreover, direct integration into the JPEG standard allows small publications and independent creators to produce verifiable content without massive technology investment. Institutional support — from the American security agencies to the international standards bodies — turns the digital signature into fundamental infrastructure of the media landscape, establishing a new paradigm where verifiable provenance is both a recommended practice and a legal requirement.

The promise and the peril

In a world where a photograph has been, for a hundred years, the most direct proof of an event — a visual testimony that speaks without need of translation — the irruption of generative artificial intelligence is rewriting the rules of the game. The promise is clear: with **Content Credentials** (C2PA) and an

immutable anchor on the blockchain, every image can carry a digital signature verifying who captured it, at what moment and under what conditions. That would restore trust in visual journalism, distinguish verifiable content from manipulative noise and allow photographers to prove unambiguously their authorship and the authenticity of their coverage.

The first steps already show the potential. In February 2025, **AFP** and **Nikon** presented a prototype camera that automatically signs each capture with C2PA, recording the provenance information at the very moment the shutter opens [4]. Since then, brands like **Leica** (with the M11-P), **Canon**, **Sony** and even **Google Pixel** have integrated secure chips that store the credentials in hardware [6]; key institutions — from the **NSA** and **CISA** to the **JPEG Trust** standards consortium — have recommended the use of these signatures as standard practice since January 2025 [8]. In this scenario, the image becomes a document as reliable as a bank transaction.

Yet the danger hiding behind this appearance of security is real and complex. The same information that guarantees authenticity can reveal sensitive data: GPS coordinates, identity metadata or even clues about the source of a piece of journalism. In conflict zones, that trail can be clue enough for repressive forces to locate and silence photographers and activists. The case of the fake "Pentagon explosion" of May 2023, analyzed by **Hany Farid**, shows how a single AI-generated image, spread by an account posing as a news outlet, briefly moved the financial markets, underlining the fragility of an ecosystem in which provenance is the only reliable filter [1].

There is also the risk of creating a **bifurcation of journalism** — a world divided between verified and unverified content — that could marginalize legitimate voices lacking advanced tools to sign their images. The "cheap fakes," which outnumbered deepfakes by a ratio of 7:1 in the **2024** elections, demonstrate that the central problem is **provenance and context**, not just the visual quality of the forgery [2]. If trust is placed exclusively in the digital seal without any independent audit mechanism, there is a danger of **overconfidence** in a certificate that can be stripped or altered; C2PA metadata are vulnerable to suppression on social platforms and in transcoding processes [7].

To preserve the promise without falling into the peril, privacy systems allowing **selective disclosure** must be designed. Photographers should have full control over which metadata are revealed and when, through encryption protocols that separate sensitive information (like GPS) from the certificate of authenticity. Media organizations could establish **anchoring policies** in which the proof of existence on the blockchain guarantees the permanence of the act, while the private metadata remain in protected environments and are released only with explicit consent.

In short, technology offers a way to trust the image as testimony again, but that trust must be balanced with rigorous data-protection measures and an ethical framework that avoids creating a two-speed journalism. Only then can we guarantee that photos remain a powerful voice for the truth, without putting at risk the lives that create them.

This anchoring architecture does not only reinforce the professional media; it creates a space where citizen journalism can become legal testimony. In Taiwan, the former digital minister **Audrey Tang** called on the population to record their own polling stations, offering a chain of video which, once anchored to the blockchain, could be instantly checked against the conspiracy theories circulating on social networks [9]. The same mechanism has found an emerging ecosystem in Peru: **Voto Libre** uses AI to classify photographic evidence of electoral fraud and registers it as immutable proof, guaranteeing "secure, transparent, real-time" reporting that can be audited without relying on intermediaries [10]. In both cases, an anchored image transcends the mere visual document: it becomes verifiable evidence of what really happened, extending trust not only to the big newsrooms but to the direct voice of the citizenry.

It's time to act

The media are the first stage on which the breakdown of visual trust becomes palpable. Every day a newsroom publishes a photo without a verifiable certificate, it leaves the door open to deception: the fake "Pentagon explosion" of May 2023, AI-generated and circulated as real news, caused a brief but significant market swing before being debunked by Hany Farid [1]. As we have seen, the first commitment to make is **adopting signed capture**: the new Leica M11-P, Nikon, Canon and other cameras already carry security chips capable of stamping a Content Credential at the moment the image is obtained [6]. Newsrooms must require that all visual material come from these devices or, failing that, be processed with a certified tool that adds an unalterable digital signature.

Furthermore, **publication must be accompanied by an immutable proof of existence**. C2PA metadata can be stripped or altered by social networks and transcoding processes, as recent critiques warn [7]. Every outlet must therefore anchor its evidence to a public blockchain recording the exact time of capture and the device's signature; this anchor guarantees that even if the credentials are erased, the reader can verify the original provenance with a simple link.

Digital platforms bear a complementary responsibility: **preserving and displaying the credentials** instead of deleting them to optimize storage or on privacy grounds. The NSA and CISA recommendations (January 2025) already indicate that content management systems must integrate native support for Content Credentials [8]; that means updating APIs, offering creators options over what information to expose and ensuring the display layer clearly shows the verification status to the public.

Camera manufacturers are the **point of origin** of this chain of trust. They must **extend capture-at-origin with configurable privacy**, letting users choose which attributes (location, personal metadata) stay encrypted and which remain visible for truth checking. The pioneering example of AFP with Nikon, which presented a prototype camera with signing built into capture in February 2025, shows the model is feasible and already being deployed in professional environments [4].

Citizens, finally, cannot remain passive. In a landscape where "cheap fakes" outnumbered traditional deepfakes by a 7:1 ratio in the 2024 elections, the central problem is the **provenance and context** of content [2]. Every user must learn to look for and interpret the indicators of credibility — the C2PA marks, the blockchain timestamps — and not stop at visual appearance. Media literacy campaigns must include practical tutorials on verifying an image with free tools available online.

Regulators, for their part, must establish **interoperable standards** that oblige media, platforms and manufacturers alike to follow common signing and anchoring protocols, avoiding an ecosystem of "two-speed journalism." Initiatives like the Content Authenticity Initiative already count more than 5,000 members (as we have seen) and must be reinforced with rules that guarantee the transparency of sources [3].

In short, in a world of synthetic images, **proving what is real becomes defending journalism**. The technology gives us the tools; it falls to each actor — media, platforms, manufacturers, citizens and regulators — to adopt them and make them visible. The moment to prove the truth has arrived: *time2act*.

---

[1] NPR — *Fake viral images of an explosion at the Pentagon were probably created by AI* (May 2023). https://www.npr.org/2023/05/22/1177590231/fake-viral-images-of-an-explosion-at-the-pentagon-were-probably-created-by-ai ↩

[2] NPR / Harvard Ash Center — *AI was everywhere in 2024's elections, but deepfakes were only part of the picture*. https://ash.harvard.edu/articles/the-apocalypse-that-wasnt-ai-was-everywhere-in-2024s-elections-but-deepfakes-and-misinformation-were-only-part-of-the-picture/ ↩

[3] Content Authenticity Initiative — *5,000 members* (2025). https://contentauthenticity.org/blog/5000-members-building-momentum-for-a-more-trustworthy-digital-world ↩

[4] Nikon Rumors — *AFP validates a C2PA certification test using a prototype Nikon camera* (Feb 2025). https://nikonrumors.com/2025/02/18/afp-validates-a-certification-test-of-its-photos-via-the-c2pa-standard-using-a-prototype-nikon-camera.aspx/ ↩

[5] IPTC — *C2PA news: France Télévisions, IBC "Stamping Your Content"*. https://iptc.org/news/tag/c2pa/ ↩

[6] Content Authenticity Initiative — *Leica launches world's first camera with Content Credentials*. https://contentauthenticity.org/blog/leica-launches-worlds-first-camera-with-content-credentials ↩

[7] arXiv — *Why the C2PA specifications fall short* (strippable metadata). https://arxiv.org/html/2604.24890v1 ↩

[8] NSA/CISA — *Content Credentials: Strengthening Multimedia Integrity* (Jan 2025); JPEG Trust ISO/IEC 21617-1. https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF ↩

[9] Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 3. ↩

[10] Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 3; Infobae (2025). https://www.infobae.com/peru/2025/06/06/blockchain-y-elecciones-2026-como-se-desarrollara-en-peru-el-voto-digital-obligatorio-para-militares-y-ciudadanos-en-el-extranjero/ ↩

Life or Death of Data: Healthcare and Medical Imaging

"As to diseases, make a habit of two things — to help, or at least to do no harm."

— Hippocrates, Epidemics, I, XI

It was a cold October afternoon, and the respiratory oncology laboratory of a large university hospital hummed with its usual routine: patients waiting for the results of their scans, doctors reviewing reports and technicians uploading DICOM images to the central system. In the reading room, a veteran radiologist with more than ten years of experience opened a new computed tomography (CT) scan of a 58-year-old man who had come to the department with a persistent cough and weight loss. The image on screen showed a 1.2 cm pulmonary nodule at the upper tip of the right lung, a small dense shadow standing out against the healthy tissue.

Without hesitation, the radiologist activated the oncology protocol: the area was marked for an image-guided biopsy and diagnostic confirmation was scheduled before any treatment. Trust in the data that had arrived through the hospital network was total; the DICOM format, a standard for decades, seemed a guarantee that the image was a faithful testimony of the patient's body.

But inside, an invisible layer of algorithms could have altered the very essence of the file. A group of researchers from Ben-Gurion University of the Negev, together with the Soroka medical center, demonstrated at USENIX Security 2019 that a generative adversarial network — called CT-GAN — can inject or remove tumors in CT and MRI images in milliseconds, leaving no trace perceptible to the human eye or to conventional verification systems. In a blind experiment, radiologists classified 99% of cases with an injected tumor as malignant and 94% of cases with a removed tumor as healthy; even when warned, the error rates remained alarmingly at 60% and 87% respectively [1].

In this scenario, the attack was no science-fiction plot but published, replicable code. Someone — perhaps an external actor with access to the hospital's internal network, or a malicious insider — could have loaded the CT-GAN model onto the imaging server and silently modified the DICOM file before it reached the radiologist's terminal. The nodule she was looking at might be nothing more than a digital illusion, an artifact generated by a neural network trained to deceive.

The impact of such a manipulation is immediate and potentially lethal: the patient would undergo unnecessary invasive procedures, exposed to radiation and to the psychological uncertainty of living with a false diagnosis. On the other face of the same monster, another image could have hidden a real tumor, delaying treatment and endangering the patient's life.

This scene is not a dystopian tale; it is a warning grounded in real evidence that lays bare the fragility of modern medicine when data become vulnerable. If we cannot trust the integrity of every pixel, the very essence of clinical practice crumbles, and health — the most fundamental right — is left at the mercy of a new form of war: the war over data.

When the medical image lies

In 2019, research presented at USENIX Security revealed a disturbing scenario for healthcare: **CT-GAN**, a conditional generative network capable of injecting or removing lung tumors in computed tomography (CT) scans and brain lesions in magnetic resonance images (MRI) in mere milliseconds. In a blind experiment, radiologists diagnosed the great majority of injected-tumor cases as malignant and judged the great majority of removed-lesion cases healthy; **as we have seen**, the clinical error was considerable.

The **DICOM** format (Digital Imaging and Communications in Medicine) defines the encoding, storage and transmission of medical images, but it predates modern cybersecurity and **incorporates no intrinsic cryptographic mechanism**. Images travel across hospital networks, VPNs or cloud services often without digital signatures that could guarantee their integrity. Into this void, an attacker with access to the supply chain — an IT infrastructure provider, for example, or a malicious insider — can introduce the CT-GAN model and modify data in real time, with no visible record of the change.

The motivations are diverse. In **insurance**, a falsified image could generate coverage for nonexistent treatments or, conversely, hide injuries to reduce payouts. In biomedical research, image manipulation

can lead to misleading clinical results, diverting resources and feeding fraudulent publications. At a more sinister level, a hostile state or group could **sabotage** critical clinical studies or provoke misdiagnoses that endanger patients' lives, turning the laboratory into a new front line of war [1].

The study's results show that neither manual review nor traditional alert systems are enough; the radiologists, even informed of possible manipulations, trusted their visual experience and failed in most cases. This exposes a structural vulnerability: **blind trust in the image as absolute truth**. In an ecosystem where every pixel can be altered in milliseconds, the cost of a wrong diagnosis translates immediately into inappropriate treatments, unnecessary interventions and, in the worst scenario, the patient's death.

To reverse this trend, *time2act* proposes **the proof of existence**: generating a cryptographic hash of every DICOM file inside the acquisition device and sealing it with a timestamping authority (TSA) under RFC 3161. The fingerprint is anchored on the public blockchain, optionally aggregated into a Merkle anchor, with that timestamp as a complement kept off-chain; this creates an immutable trail that complies with HIPAA and 21 CFR Part 11 [5]. Any later alteration would be detected by comparing the original hash with the new value, opening the door to automatic audits and an early-warning system.

As we have seen, the CT-GAN example illustrates how a manipulated data point can turn a medical image into a deadly weapon. The answer is not only technological but regulatory and cultural: establishing **digital chain-of-custody** protocols, requiring cryptographic signatures for every study and educating clinical staff about the fragility of traditional formats. Only then can we guarantee that images remain faithful testimonies of health, and not forgeries that gamble with human lives.

Synthetic doctors and patients

Imagine a doctor who, first thing in the morning, receives a call from a patient complaining of unbearable abdominal pain; the voice is clear and the communication style familiar — she even mentions an anecdote from the last visit. Suspecting nothing, the doctor authorizes a prescription for strong painkillers and orders an urgent ultrasound. A few hours later, another professional notices that the social security number on the record matches no existing file: the "patient" was a synthetic identity generated by audio and video deepfake, run by a criminal network dedicated to obtaining prescriptions for the black market [2].

This scenario is no anecdotal exception; it represents an emerging threat that goes beyond simple pharmaceutical fraud. Voice- and image-generation technologies — based on deep-learning models such as generative adversarial networks (GANs) — can replicate with surprising fidelity the face and tone of a well-known doctor or a regular patient. When these imitations spread through telemedicine platforms, the traditional security layer — firewalls, TLS encryption and multi-factor authentication — proves insufficient: none of these mechanisms "sees" a fabricated face or detects a cloned voice. The danger lies in human perception; professionals trust the authority conveyed by facial expressions and vocal timbre, and malicious networks exploit that trust to manipulate clinical decisions [2].

Moreover, the phenomenon of "synthetic patients" feeds a chain of fraud that begins in the teleconsultation and stretches into the insurance systems. Organized groups can create fake profiles with digital identity documents, simulated medical histories and doctored radiological images; these data are used to file inflated compensation claims or to obtain reimbursement for treatments that never took place. In some cases, a cloned doctor's voice tricks staff into releasing confidential data or signing authorizations for nonexistent procedures, causing immense economic cost and loss of confidence in the health system, especially amid the mass adoption of telemedicine [2].

The solution is not merely to harden the technical barriers, but to institute a digital chain of custody based on cryptographic proofs of existence. Every medical image, voice recording and clinical document should be associated with a unique hash sealed by a timestamping authority (TSA), as we have seen with RFC 3161, ensuring the file never leaves the original device without a verifiable mark. This approach complies with privacy rules, including HIPAA and 21 CFR Part 11, and creates an immutable trail that defuses any attempt at post-capture manipulation [5][7].

Ultimately, technology and culture must combine: deploying blockchain solutions for storing minimal metadata off-chain, training clinical staff to detect voice and image anomalies, and establishing multi-factor authentication protocols that include biometric checks independent of the data presented. Only

then can we guarantee that the trust placed in a virtual consultation does not become a weapon for those who would gamble with human life.

The integrity of science

In an era when biomedical journals are the epicenter of clinical knowledge, contaminated data are as deadly as a tumor injected into a CT scan. The image detective Elisabeth Bik has shown that close to **3.8%** of the more than **20,000** biomedical articles she reviewed contain problematic figures — odds of one in every 25 papers — and that half of those cases show clear signs of deliberate manipulation [3]. That rate translates into a systematic erosion of the foundation on which therapeutic decisions are taken, protocols designed and public resources allocated.

The neuroscientist Hoau-Yan Wang, linked to the development of Cassava Sciences' Alzheimer's drug, was federally indicted on **June 28, 2024** for fraud, including the fabrication of Western blots supporting false results [3]. The investigation revealed that approximately **$16 million** in funding had been obtained from the National Institutes of Health (NIH) on the basis of manipulated data, diverting resources meant for treatments for millions of patients.

The rise of automated platforms like Proofig and ImageTwin, already integrated into editorial workflows, offers concrete hope [3]. These tools apply detection algorithms based on statistical pixel analysis, compression patterns and comparisons against reference libraries, identifying anomalies a human review might miss; their effectiveness depends on universal adoption and an institutional commitment to validate every figure before publication.

The **danger** is twofold. First, bad science feeds wrong clinical decisions: doctors who trust false results may prescribe ineffective or even harmful treatments, putting human lives at risk.

Second, public funding is squandered on projects built on unreliable data, which shrinks the resources available for genuine research and deepens citizens' distrust of the health system.

The promise lies in an integrity architecture based on proofs of existence. Using hash + timestamp (TSA) under RFC 3161, images and metadata are recorded off-chain but with an immutable fingerprint on the blockchain, meeting the privacy requirements of HIPAA and 21 CFR Part 11 [5][7]. Each file preserves its integrity inside the clinical device, generating an audit trail that any auditor can verify in real time.

The message, then, is clear: without a verifiable layer of trust, science becomes a minefield. As we have seen, the academic community and the regulatory institutions must systematically implement these proofs of existence and use tools like Proofig or ImageTwin to filter every figure before it reaches the reader; only then will the knowledge generated be a solid pillar for human health, not a time bomb threatening lives and funding.

Sealing clinical truth

In a world where merely injecting or erasing a pixel can turn a diagnosis into a sentence, the solution must be as rigorous as the rules that govern health itself. The emerging pattern — already endorsed by the FDA — is to keep the sensitive data entirely off the blockchain and to record only its cryptographic digest (hash) together with minimal metadata: timestamp, user identifier and type of operation. This small digital fingerprint becomes an immutable trail that can be audited without revealing confidential content, meeting or even exceeding the audit requirements of HIPAA and 21 CFR Part 11 [5].

The mechanics are simple but powerful. Every DICOM image, every clinical record and every trial file is processed inside the device or the hospital server; a SHA-256 hash representing its complete integrity is attached. That hash is transmitted to a blockchain where it is registered with a timestamping authority (RFC 3161) or, in more advanced environments, with the Merkle root of a private health chain [7]. The associated metadata — only what is needed to validate the transaction — remain visible on the chain, while the original content never leaves the secure perimeter of the medical center, thus preserving the privacy HIPAA demands and enabling transparent auditing: any review can compare the stored hash with the current file and detect, in milliseconds, any alteration.

In clinical trials this model creates a verifiable chain of custody recording who generated each data point and when it was modified. The Triall product already offers this functionality: every trial data point is "sealed" with its hash and linked to an immutable record, ensuring that regulators and sponsors can trace

the information's trajectory from the moment of collection to final publication [6]. As we have seen, image fraud affects some 3.8% of papers and has produced indictments over some $16 million in NIH funding [3].

The FDA has recognized this need. Its final cybersecurity guidance for medical devices, published on **June 27, 2025**, requires manufacturers to include a software bill of materials (SBOM) and exhaustive documentation of protection measures [4]. Any device that emits radiological images or monitoring data must incorporate mechanisms demonstrating that the content has not been manipulated after acquisition.

In practice, only three steps are needed: (1) generate and store locally the hash of each sensitive data item; (2) send that hash to a blockchain with a certified timestamp; (3) preserve the metadata for future audits. Tools like Proofig or ImageTwin automatically detect anomalies before figures reach publication and integrate into the hospital workflow. When these solutions are deployed systematically, trust in medical images shifts from an exploitable vulnerability to a verifiable guarantee.

The result is clear: an ecosystem in which clinical truth is "sealed" and cannot be altered without leaving a trace. In a context where every diagnosis can decide between life and death, this verification layer goes beyond the regulations and offers a real defense against the digital deception threatening global health. The moment has come to adopt it as a universal standard and turn the security of medical data into an immovable pillar of the future of healthcare.

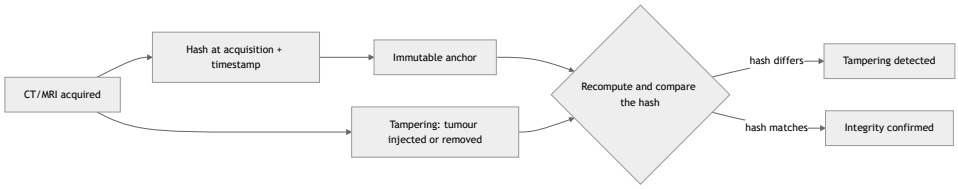


Figure 18.1 — If the image's hash is computed at the moment of acquisition and anchored with a timestamp, any later manipulation (an injected or erased tumor) is detected instantly.

How it works, without the data leaving

At the core of a trust solution for medical data lies an operation apparently simple yet extraordinarily powerful: computing a cryptographic hash of the file at the exact moment of acquisition. This mathematical digest — a unique string of characters representing the entire content — becomes the digital fingerprint of the image, the radiology report or the clinical record. Once generated, the hash is sent to a Timestamp Authority (TSA) under the protocol we have already cited, which issues a time certificate signed with its private key, attesting that the fingerprint existed at that instant and has not been altered since.

To amplify tamper resistance and enable open verification, the certified hash can be anchored to a public chain via a Merkle root. In this scenario, multiple files generate child nodes grouped into a binary tree; each node contains the combined hash of its children, and the top layer (the root) is a compact representation of the entire dataset. Publishing this root on a blockchain — with the content kept off-chain, leaving only the fingerprint and minimal metadata — creates an immutable audit trail that meets the privacy requirements established by HIPAA and 21 CFR Part 11 [5].

The key point that makes this mechanism viable for healthcare is that the original file **never leaves the device**. Only the fingerprint — a byte string of fixed length — travels to the outside environment, whether to the TSA or to the blockchain. This removes any need to transfer complete radiological images, which contain sensitive personal information, to third-party servers. Subsequent verification is fully decentralized: any interested party can recompute the file's hash locally and compare it with the fingerprint recorded in the certificate or on the public chain; if they match, the content is proven to have suffered no alteration since the moment of capture.

This architecture offers a double layer of security. On one hand it protects patient confidentiality, because the clinical data remain within the controlled perimeters of the medical device or hospital system; on the other it provides a verifiable proof of existence for auditors, regulators and health professionals without depending on any central entity that could deny access or manipulate the records. The "t2p" model applied to healthcare turns every clinical study, every diagnostic image and every digital prescription into an immutable testimony of its own origin.

In practice, this flow can be integrated into existing DICOM systems: when the scanner generates a new CT, the management software immediately computes the hash, requests the timestamp and updates the file's metadata with a link to the blockchain transaction. The professionals reviewing the images can see, next to the report, a QR code or a hash reference that lets them verify the data's authenticity in real time. Any attempt to inject false tumors — as the CT-GAN experiments demonstrate, where radiologists were deceived in up to 99% of cases — thus becomes ineffective: without the original hash and its timestamp, the manipulation could not slip past an independent auditor [1] (as we have seen).

The mechanics of the proof of existence combine three fundamental principles: **cryptographic integrity**, **certified temporality** and **preservation of data ownership**. These characteristics ensure that clinical information remains both secure and verifiable, turning it into a robust defense against the digital deception that threatens patients' lives.

The promise and the peril

Imagine a radiologist who, unknowingly, reviews a CT scan in which a tumor has been artificially inserted by a generative AI like CT-GAN. In the USENIX Security 2019 study, professionals diagnosed malignancy in **99% of cases** with injected tumors and judged healthy **94%** of the studies where the tumor had been removed; even when warned of the possible deception, the error rate stayed between **60% and 87%** [1]. These figures show that blind trust in an image can translate directly into wrong treatment decisions, with potentially lethal consequences.

The **promise** of the cryptographic seal — an immutable hash and a certified timestamp (TSA or Merkle anchor) that never leaves the device — is to provide a layer of authenticity that becomes the cornerstone of reliable diagnosis. When every DICOM study comes with a verifiable hash, the radiologist can confirm that the image has not been altered since the moment of acquisition. This not only protects the patient against malicious manipulation; it also strengthens **research**: biomedical papers containing problematic figures account for roughly **3.8%** of the studies reviewed, according to Bik's analysis [3]; an audit based on proofs of existence could catch these irregularities before publication, cutting the costs of retractions and raising the field's credibility.

In **clinical trials**, platforms like Triall offer chains of custody that record every step of the process, from sample collection to final analysis, with a hash and a timestamp guaranteeing that no data point can be deleted or modified without leaving a trace [6]. This allows regulators — for instance the FDA, which on June 27, 2025 published definitive guidance on medical-device cybersecurity requiring SBOMs and exhaustive documentation [4] — to verify the integrity of every element of the trial, preventing frauds like the fabricated Western blots case, tied to some **$16 million** from the NIH [3].

The **danger**, however, lies in a false sense of immunity. A cryptographic seal is no substitute for hospital network cybersecurity; if the source of the data (the device or the server storing it) is compromised, the hash may be generated from information already altered. The complexity of integrating this layer into clinical workflows — often saturated with legacy systems and rigid protocols — can produce delays or operational errors without clear governance over who is authorized to verify and validate the seals. Overreliance on the proof of existence could lead to **neglect** of the other attack vectors: vulnerabilities in internal networks, malware that manipulates data before signing, or the misuse of cryptographic keys by unauthorized staff.

As we have seen, the timestamp and the hash constitute a necessary layer within an integral strategy of security and integrity. When paired with robust network controls, transparent governance policies and continuous training of health personnel, they make it possible to restore trust in data-driven medicine — a trust that can mark the difference between life and death for millions of patients.

It's time to act

The integrity of medical data is no longer a technical question; it is a vital requirement that every entity in the health sector must adopt as an operating norm. The first responsibility therefore falls on **hospitals and equipment manufacturers**: the moment a CT scan or a medical record enters the system must be permanently registered with a cryptographic hash and anchored to a timestamp (TSA) under RFC 3161. This process, completed in milliseconds, prevents any later manipulation from passing without a trace of inconsistency, as the CT-GAN case demonstrates, where the insertion or removal of a tumor is immediately detectable when a reliable hash check is applied [1].

Researchers and scientific journals must make the timestamp a condition for accepting any article containing biomedical images. Tools like Proofig or ImageTwin already offer integrity-audit services; their systematic use should be as routine as peer review. That would drastically reduce the impact of frauds like those detected by Bik, where close to 3.8% of papers presented problematic figures [3], and would help prevent cases like the federal indictment of June 28, 2024 over fabricated Western blots, tied to some $16 million from the NIH [3].

Regulators, led by the FDA with its final cybersecurity guidance (June 27, 2025) and the EMA, must explicitly incorporate proof-of-existence requirements into medical-device regulations. That means demanding a complete SBOM and documentation demonstrating the immutable traceability of clinical files, as recommended by the blockchain pattern for healthcare, which combines off-chain storage with hashes and minimal metadata on-chain, in compliance with HIPAA and 21 CFR Part 11 [4][5].

Insurers can turn this integrity layer into a coverage criterion: before authorizing payment for a treatment or surgery, they can validate that the underlying data have not been altered. An automated verification system based on the chain of custody (like the Triall product) could reduce the costs associated with fraudulent claims and increase payers' confidence in the authenticity of the diagnosis [6].

Finally, **patients and professionals** must demand transparency. Every professional should be able to present a digital certificate proving that the images and reports they use are exact and have not been modified since the moment of capture. Patients, for their part, can request access to this immutable record as a right to information, reinforcing a culture of shared responsibility.

In short, protecting the integrity of data is protecting lives. Proof-of-existence technology is already available and validated in real scenarios; what is missing is widespread adoption and regulatory enforcement. We cannot allow a manipulated image to be the cause of a wrong diagnosis or a fraudulent prescription. As we have seen, it is time2act: let us implement hash and timestamp at every point of entry, make it mandatory for research and demand verifiable traceability at every level of the health system. Every second we wait is a potential risk to patients' lives.

---

[1] Mirsky et al. — *CT-GAN: Malicious Tampering of 3D Medical Imagery* (USENIX Security 2019). https://www.usenix.org/conference/usenixsecurity19/presentation/mirsky · summary with figures: https://www.pindrop.com/article/threat-deepfakes-in-healthcare/ ↩

[2] TruthScan / Pindrop — *AI-driven fraud in healthcare: deepfake doctors and synthetic patients* (2025). https://truthscan.com/blog/ai-driven-fraud-in-global-healthcare-2025-trends-and-countermeasures/ ↩

[3] STAT News — *Cassava Sciences scientist indicted (fabricated Western Blots)* (28-Jun-2024); Retraction Watch / Bik. https://www.statnews.com/2024/06/28/cassava-sciences-indictment-alzheimers-hoau-yan-wang/ ↩

[4] FDA — *Cybersecurity in Medical Devices: final guidance* (Federal Register, 27-Jun-2025). https://www.govinfo.gov/content/pkg/FR-2025-06-27/html/2025-11669.htm ↩

[5] HIPAA Vault — *Blockchain integration for healthcare records (audit trails)* (2025). https://www.hipaavault.com/resources/blockchain-integration-healthcare-records/ ↩

[6] Bentham/RRCT (PubMed) — *Blockchain as a Prime Guardian of clinical-trial data* (2025); Triall. https://pubmed.ncbi.nlm.nih.gov/39851126/ ↩

[7] ProofStamper — *RFC 3161 + blockchain anchoring (proof of existence).* https://proofstamper.com/en/compare/file-timestamp-tools-2026/ ↩

Credentials and Authorship: Education and Research

"I neither know nor think that I know."

— Plato, Apology of Socrates (21d)

In the cold light of an intensive care unit in south Texas, the monitor clock read 03:17 when a woman entered the room with a confident step and an arm full of paperwork. She wore an immaculate white coat, and on her chest gleamed a badge that shone as if made of gold: "RN — Registered Nurse." The patient, a fifty-seven-year-old man with severe pneumonia, watched her hand while the ventilator kept up its constant noise. The nurse leaned over the bed, checked his vitals and, without hesitation, administered a dose of the antibiotics the doctor had prescribed.

The moment seemed routine; yet the truth hidden by that uniform was anything but ordinary. That "Registered Nurse" title came from no accredited university or state-recognized program of study. It was one of the more than 7,600 fake credentials detected in Operation "Nightingale," a US Department of Justice investigation that dismantled an international network dedicated to the manufacture and mass sale of fraudulent nursing diplomas (DOJ 2023; phase II, 2025) [3]. Each credential sold for up to some $15,000, and the buyers were students and professionals already working in the field alike.

The atmosphere turned tense — the patient was showing signs of an adverse reaction to the antibiotic. The nurse, with no verifiable record confirming her training or the date her title was obtained, had been making clinical decisions on the strength of a legal appearance alone. Her "diploma" was just a digital file with a hash sealed on some unknown blockchain; the real document, if one existed at all, was not on the chain but off it, as in the MIT Blockcerts model, where the on-chain trace of a degree is limited to a hash and a sealed transaction [5].

This incident embodies the thread that ties this chapter to the previous one: the concrete danger that credential forgery can pose to human life. In chapter 18 we showed how deepfakes and identity impersonation in remote examinations can compromise academic integrity; here, the same vulnerability translates into a direct threat to health. The question that emerges is inevitable: **how can we know with certainty that a degree — or a piece of work — is authentic?**

In a world where generative AI can produce convincing-looking documents and forgery networks operate at global scale, traditional trust in signed paper or a registration number is no longer enough. The answer must turn to systems that guarantee origin and integrity without depending on intermediaries: verifiable credentials based on blockchain and proofs of authorship which, as the W3C Verifiable Credentials Data Model 2.0 (May 2025) demonstrates [6], offer a cryptographic layer of trust that can be independently audited.

While the nurse continued her care, the monitor showed an increasingly irregular rhythm. A patient's life could depend on the capacity of the educational and professional system to prove, beyond doubt, who is genuinely qualified to heal him. The silence that settled over the room was not just the sound of the machines; it was the echo of a question that should resound in every classroom, laboratory and office: **what do we do when trust disappears?**

Assessment in crisis

In the spring of 2025, the classrooms of the United Kingdom unexpectedly became a large-scale experimental chamber: among the library rows and study halls, the migration to artificial intelligence platforms was skyrocketing. That year's HEPI survey revealed that **92% of UK students** were already using AI tools to write essays, solve math problems or even generate answers to online quizzes. More alarming still, the share of students using AI **during assessments** jumped from **53% to 88% in a single year**, an evolution exceeding any historical increase in assistive technologies within higher education.

At first glance, these figures might suggest a democratization of knowledge: more students accessing advanced resources and improving the quality of their work. But the picture blurs when the limits of current detection tools come into view. The Australian regulator TEQSA concluded in 2024 that quantitative methods for identifying AI use are "unreliable." An accused student can simply deny outright

having used a generative model, with no verifiable trace left behind; at the same time, **skilled users are rarely caught**, since detection solutions rely on superficial patterns rather than cryptographic proof.

This vulnerability has tangible consequences. TEQSA points to a growing risk of **false positives**: innocent students unjustly accused of plagiarism or of resorting to AI, which can translate into academic sanctions and damage their professional trajectories. The social cost of these errors is evident when one considers that roughly **15.7% of students** already suffered some form of contract cheating before the mass arrival of generative AI, a figure amplified by the absence of reliable evidence [2].

In this context, the answer cannot simply be harsher disciplinary policies. The assessment reform TEQSA proposes centers on **redefining the criteria of verification**, moving from a model based on assumptions to an ecosystem in which the proof of authorship and the integrity of documents are indisputable. That means adopting technologies like verifiable credentials stored on blockchain, which allow both the issuing of a degree and the exact moment a piece of academic work was created to be immutably recorded.

The transition to this new paradigm does not remove the need for an ethical culture; on the contrary, it provides a solid technical basis for institutions to **distinguish clearly** between a student's genuine contribution and the support of external tools. Instead of focusing on pursuit and punishment, educational institutions can channel resources into systems that guarantee the transparency of the assessment process, reducing both false positives and the opportunities for mass fraud.

In short, the current landscape shows a **deep erosion of trust** in academic assessment, driven by mass AI adoption and by the inability of traditional detection mechanisms. The lasting solution runs through an infrastructure of verifiable credentials that restores the security of the degree and of authorship, laying the foundations of an educational environment resilient to emerging technologies.

Degrees that are worthless

Imagine an official document bearing the seal of a nursing academy: glossy paper, a registration number and an apparently authentic digital signature. A student takes pride in having earned that degree after an intensive online course, but no one sees that the diploma belongs to a clandestine network like the one uncovered by Operation Nightingale, an investigation led by the US Department of Justice.

As we have seen, in 2023 more than **7,600** fake diplomas and transcripts were discovered in Florida; the first wave alone brought **25 defendants** and the revocation of licenses in several states, and left patients exposed to unqualified professionals. Phase II, opened in 2025, added twelve more defendants and confirmed that forgery remains a growing threat [3]. While the authorities fight this underground market, **deepfakes** — AI-generated images that superimpose one person's face and voice onto another's body — are spreading into the world of remote exams and interviews. A student can now pass off their identity before an exam camera with a believability that fools algorithms and human proctors alike, and these manipulations can be stored and reused for future assessment opportunities, destroying trust built over decades [4].

The common pattern between the two threats is clear: **credentials and identity can both be forged**. The degree is no longer a reliable seal of competence; the face on the screen does not guarantee the candidate's real presence — exposing a fundamental vulnerability in the current educational model, based on printed documents or digital signatures that are easily replicated and manipulated.

The solution is not just stronger detection: what is needed is a **verification infrastructure that does not depend on eyeballing**. Blockchain-based verifiable credentials encode every degree and identity record with a unique hash and seal them on an immutable chain, making modification impossible without leaving a trace. This model guarantees the document's authenticity and lets any interested party (universities, regulators, hospitals) check its validity instantly, without intermediaries.

In this context, credential fraud and digital impersonation are an urgent call to action: decentralized systems must be adopted to restore trust in education and health, so that degrees are truly **degrees that count**.

The answer: verifiable credentials

The first step toward an education that recovers trust was already tested in MIT's laboratory, where the Blockcerts project ran a pilot of **111 diplomas issued on blockchain** in 2017. In this scheme the degree's sensitive data — name, date of award and program of study — stay off-chain; all that is embedded in the network is a **one-way hash** representing the document's full content and a sealed transaction certifying that MIT created the record. This architecture is identical to the hash-anchor model used by the proof of authorship/existence, where the same mechanism demonstrates that someone possessed a file at a given moment and that it has not been altered since [5].

In Europe, the approach has been scaled to continental level thanks to **EBSI (European Blockchain Services Infrastructure)** and the framework of the **European Digital Credentials for Learning**. These platforms issue credentials with **eIDAS electronic seals**, ensuring legal validity and cross-border interoperability. More than **76,000 verifiable credentials** have been generated so far, and any company or regulatory body can consult them instantly, with no intermediaries or manual processes. Standardization arrived with the **W3C Verifiable Credentials Data Model 2.0 Recommendation** (May 2025), which defines how to package, sign and validate these credentials securely and compatibly with any public blockchain infrastructure [6].

The added value of this model is double: it eliminates the risk of **fake diplomas**, like those detected in Operation Nightingale, where more than **7,600 forged nursing certificates** endangered patient safety in the US (DOJ 2023; phase II with 12 additional defendants in 2025) [3]; and it counters **deepfake impersonation** and other forms of digital proctoring abuse that have compromised remote exams, since each student answer can be tied to a verifiable credential guaranteeing the origin and integrity of the work [4].

In practice, an employer or a regulatory authority only has to enter the degree's **hash** into the corresponding verification interface; the blockchain returns the timestamp and the provider's digital signature (MIT, a European university, etc.), confirming that the document has not been modified since issuance. No calls to the institution, no emails, no weeks of waiting for an official response: verification completes in seconds and is **immutable**. This creates an ecosystem in which a degree's value ceases to hang on the reputation of the issuing entity — which can fall prey to fraud or administrative failure — and passes directly to the **cryptographic record**, which is transparent, auditable and tamper-resistant.

According to the 2025 HEPI report, as we have seen, AI use in assessment has surged; with blockchain-based verifiable credentials, universities can erect a technical barrier that prevents AI from producing forged documents without a trace, reinforcing trust in issuer and recipient alike.

In the end, the combination of **blockchain diplomas**, **eIDAS seals** and the **W3C Verifiable Credentials 2.0** standard constitutes a concrete, proven answer that restores the two basic forms of trust in education: recognition of the degree and proof of authorship. Now more than ever is the moment to adopt these decentralized technologies and guarantee that our degrees are **truly degrees that count** [1].

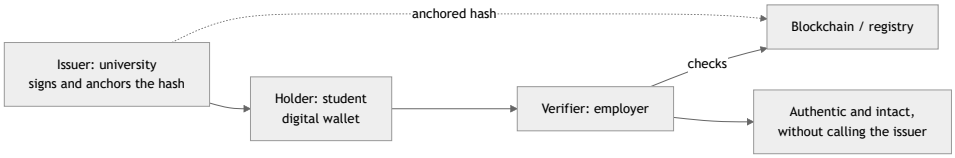


Figure 19.1 — The issuer signs and anchors the credential's hash; the holder presents it; any verifier checks its authenticity and integrity without intermediaries.

Proving authorship

In a world where generative AI can reproduce text with a fidelity that blurs the lines between original and copy, the true proof of authorship becomes the cornerstone of any reliable educational system. The solution does not lie in trying to prove the absence of intervention — a practically impossible task when the algorithms grow ever more opaque — but in certifying **WHEN** and **WHAT** was created. That is the premise of the proof of existence: a unique hash, combined with an immutable timestamp, demonstrates that the creator possessed the document at a given moment, establishing the priority of authorship before the content reaches any review or publication.

The mechanism is simple but potent. Generating the hash — a fixed-length cryptographic digest — captures the complete state of the file at that instant and records it on a public blockchain, where each block carries a timestamp verifiable by the decentralized network. The result is irrefutable proof that the

document existed and had not been altered since that moment [7]. This offers two crucial advantages: first, against plagiarism, the creator can demonstrate "I had it first"; and second, against accusations arising from unreliable detectors' false positives, there is objective evidence that dispels any doubt about the work's legitimacy.

As we have seen, the MIT/Blockcerts project certified 111 studies in 2017 and the European Union has stored more than 76,000 credentials as of May 2025, under the W3C Verifiable Credentials 2.0 model, proving that the infrastructure for proof of authorship is already operational and scalable [5][6].

The same principle applies to the intellectual work of researchers and students. Every article, every thesis project, even every piece of source code can be "sealed" before being shared with supervisors or reviewers. This not only protects the author against unjust accusations; it also strengthens confidence in academic evaluation processes; as we have seen, AI use in assessment has recently exploded [1]. In this context, institutions cannot keep depending solely on authorship detectors that are "unreliable" and that often leave room for denial [2].

Tangible proof of this practice can be found in the very work you are reading. Before publication, the manuscript of **time2act** was hashed and sealed on a blockchain, guaranteeing that no one — not even the author — could modify it without leaving a trace. This gesture is no mere technical exercise; it is the concrete manifestation of the principle this chapter defends: **trust is rebuilt with verifiable proof, not with assumptions**.

The call to action is clear. Educators, administrators and researchers must immediately adopt blockchain-based proof-of-existence tools for all their intellectual products. That means integrating solutions like Blockcerts or EBSI into academic workflows, establishing policies that mandate pre-publication sealing and training staff in the management of these digital credentials. Only then can we guarantee that degrees and intellectual production are truly **degrees that count**, and that our education recovers its two fundamental trusts: recognition of the degree and proof of authorship.

The promise and the peril

Imagine the digital classroom of a European university where a professor has just published the final results of a bioengineering course. Once the file is uploaded to the institutional repository, an algorithm automatically generates a *Verifiable Credential* based on the W3C 2.0 model (recommended by the W3C in May 2025) and names it "Bioengineering Diploma." The student who passed the course instantly receives a blockchain credential with a unique hash and a verifiable timestamp. With no paper and no phone calls to the central registry, any company or authority can confirm the degree's validity in seconds: a simple query to the public EBSI explorer or the corresponding national node.

This fluidity is the promise underpinning the whole paradigm of digital credentials. A degree becomes a portable, interoperable, immutable asset; the proof of authorship, through the same hash + timestamp mechanism, guarantees that the student owned the work at the exact moment of its creation (as the MIT/Blockcerts project demonstrated with 111 students piloted in 2017) [5]. Verification is instant and decentralized: companies can validate thousands of job applications without depending on a central certification office, cutting administrative costs and eliminating frauds like the 7,600 fake diplomas detected in Operation Nightingale in the US (2023-2025) [3].

Yet the same infrastructure that provides trust can generate new dangers. Invasive *proctoring*, driven by AI-based monitoring platforms and facial recognition, collects biometric data from students under 18 without a clear framework of consent or control over the subsequent use of those records. Privacy comes under threat when institutions store images, audio and behavioral patterns on centralized servers or, worse, on public chains without adequate anonymization.

Academic fairness, moreover, risks becoming a new digital divide. According to the HEPI 2025 study, 92% of UK students already use AI, and the share using it in assessments jumped from 53% to 88% in a single year [1]. Whoever lacks access to devices capable of generating or verifying blockchain credentials will be relegated to slower, less recognized manual processes, widening the inequality between well-funded institutions and those operating on limited resources.

Dependence on a single sealing-service provider is also a structural vulnerability. If a company controlling key infrastructure suffers an attack or decides to change its terms of service, the whole educational chain can see the integrity of its degrees and proofs of authorship compromised. The risks of false positives —

and the 15.7% baseline of contract cheating detected in Newton's study — demonstrate that no system is infallible [2].

To seize the promise without falling into the perils, an architecture centered on privacy and equity must be designed. Selective disclosure must let the student decide what data to share with third parties; control over their credentials must be built into friendly interfaces that require no advanced technical knowledge. Institutional policies, moreover, must include decentralized verification mechanisms and independent audits to guarantee that providers maintain robust security standards.

In short, blockchain can turn education into an ecosystem where reliable degrees and fair authorship are the norm; but only if personal data are protected, equal access is achieved and the trust infrastructure is diversified. Adopting these technologies now will decide whether our educational systems advance toward genuine *trust-by-design* or become weak points for abuse and discrimination.

It's time to act

Universities and institutions of higher learning are the first front where the crisis of trust translates into concrete losses: fake diplomas reaching operating rooms ([3]), exams impersonated by deepfakes threatening the integrity of selection processes ([4]) and a generation of students of whom, as we have seen, 92% already use AI in their academic work, with a 35-percentage-point jump in assessment use in a single year ([1]). The solution does not lie in relying on ever more sophisticated detectors — which, as TEQSA warns, are "unreliable" and often leave the accused free to deny use ([2]) — but in rebuilding the credentialing model from its core.

1. Universities and institutions: issue verifiable credentials

Adopting open standards like the W3C Verifiable Credentials Data Model 2.0 (May 2025) ([6]) turns every academic degree into a "hash + sealed transaction" on the blockchain, as the MIT Blockcerts project demonstrated with 111 students piloted in 2017 ([5]). These credentials are immutable, instantly verifiable and compatible with the European digital identity systems (EBSI) that already manage more than 76,000 certificates ([6]). Universities must build their own infrastructure or join federated networks to issue and register these seals, while guaranteeing student privacy through selective disclosure of data.

2. Students and researchers: seal the proof of existence of your own work

Every article, project or thesis must be timestamped on a public blockchain before it circulates. The same mechanism that protects this book (hash + timestamp) secures the priority of authorship and prevents impersonation ([7]). Research management platforms must integrate APIs that automate this process, giving the author verifiable proof that any reader or reviewer can consult.

3. Employers and professional bodies: verify credentials instantly

Hiring and accreditation processes must incorporate real-time queries against the blockchain registry of degrees, removing the need for phone calls or scanned documents that can be forged. This reduces the risk of hiring professionals with fraudulent diplomas — a problem already laid bare by Operation Nightingale ([3]) — and increases confidence in declared competencies.

4. Regulators and accreditation agencies: recognize interoperable standards

Regulatory bodies must formalize official recognition of verifiable credentials, establishing compliance frameworks compatible with eIDAS and other international rules. This creates a layer of trust-by-design at macro scale, avoiding the fragmentation of the educational market and ensuring that the new technologies do not become weak points for abuse or discrimination.

5. Educational platforms: integrate provenance with privacy

Online learning providers must design systems in which the intellectual property of content and users' credentials are managed with selective disclosure, as the verifiable-data model recommends. This lets students demonstrate competencies without revealing unnecessary personal information, aligning privacy with transparency.

Protecting the authenticity of the degree and of authorship is, in the end, defending the very merit of education. Every institution that adopts verifiable credentials and proofs of existence not only strengthens

the trust of its graduates; it contributes to a global ecosystem in which academic truth stays intact against the threats of generative AI. It is **time2act**: the future of learning depends on our capacity to build it today.

---

[1] HEPI — *Student Generative AI Survey 2025*. https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/ ↩

[2] TEQSA — *The evolving risk to academic integrity posed by generative AI* (2024). https://www.teqsa.gov.au/sites/default/files/2024-08/evolving-risk-to-academic-integrity-posed-by-generative-artificial-intelligence.pdf ↩

[3] U.S. DOJ — *Operation Nightingale: fraudulent nursing diplomas* (2023; Phase II 2025). https://www.justice.gov/usao-sdfl/pr/fraudulent-nursing-diploma-scheme-leads-federal-convictions ↩

[4] Open edX — *Proctoring in the age of AI: deepfake impersonation* (2025). https://openedx.org/blog/proctoring-in-the-age-of-ai-tackling-modern-cheating-techniques/ ↩

[5] MIT Technology Review — *Digital diplomas (Blockcerts)* (2018); MIT Registrar. https://www.technologyreview.com/2018/04/25/143311/digital-diplomas/ ↩

[6] European Commission — *EBSI Verifiable Credentials*; Europass *European Digital Credentials for Learning*; W3C VC Data Model 2.0. https://europass.europa.eu/en/european-digital-credentials-learning ↩

[7] ProofStamper / OriginStamp — *Proof of authorship via hash + timestamp*. https://proofstamper.com/en/compare/file-timestamp-tools-2026/ ↩

Truth as a Weapon: Defense and Human Rights

"I swore never to be silent whenever and wherever human beings endure suffering and humiliation."

— Elie Wiesel, Nobel Peace Prize address (1986)

It was March 16, 2022, three weeks after Russian forces broke through the Ukrainian border and launched an offensive that shook the whole of Europe. At dawn, the mobile phones of diplomatic offices, the tablets of war correspondents and the screens of half a country showed the same message: a video in which President Volodymyr Zelensky appeared to ask his soldiers to "lay down their arms." The footage, broadcast by the channel Ukraine 24 — which had been hacked that morning — was projected onto the station's news ticker and, simultaneously, embedded on its website.

The clip was evidently a forgery: low resolution, blurred facial contours and a skin tone that did not match the president's natural color. Yet in the frantic information race of an active war, the video spread across social networks like a wave of disinformation, feeding rumors and raising a first alarm: no leader of the country was calling for surrender — and perhaps there was an internal rupture to exploit.

Within a few hours, however, the official response was as swift as it was emphatic. Zelensky published a video — an agile, authentic moment showing his face without filters — and, in a clear voice, flatly rejected the demand: "We are not going to lay down any weapons." The message, distributed through the Presidency's verified channels and relayed by the international agencies, defused the forgery within a day.

That episode stands today as the first documented case of a "weaponized deepfake" — an audiovisual forgery deployed directly as a psychological weapon inside an ongoing armed conflict. Its impact was not limited to momentary confusion; it demonstrated that whoever controls the perception of truth can gain a decisive strategic advantage, altering military decisions and the people's confidence. The information war, an invisible but vital front, shows that truth itself becomes a scarce and contested resource.

In this context, the challenge is not merely to detect deepfakes once they have circulated; a durable defense requires authenticating communications from the moment of their creation and sealing them with technologies like cryptographic hashing and blockchain, ensuring that every official statement, every eyewitness image, carries unalterable proof of origin. That need becomes the core of our proposal: to turn truth into a defensive weapon for the protection of human rights and international justice. [1]

The weapon of disinformation

Over the past decade, influence operations have moved beyond simple propaganda delivery to become a digital architecture that manipulates realities and shapes political decisions at global scale. The Russian campaign known as "Doppelganger," for example, has deployed a network of bots and artificial intelligence systems capable of cloning the websites of Western media and governments, distorting news and amplifying narratives hostile to Ukraine. These forged platforms appear in search engines as legitimate results, deceiving casual readers and professional analysts alike. The European Union, recognizing the gravity of the phenomenon, has imposed sanctions on the entities linked to the project, stressing that its aim is not merely to disinform but to erode confidence in democratic institutions [2].

2024 marked a turning point: a global election season coinciding with the explosion of deepfakes across several continents. In Slovakia, days before the vote, a fake audio recording attributed to a political leader circulated, distorting the official discourse and stoking social discontent. In the United States, an automated call — a "robocall" — played the cloned voice of President Joe Biden explicitly asking citizens not to cast their votes. The Federal Communications Commission (FCC) proposed a $6 million fine against the political operative responsible, Steve Kramer, for this attempt at electoral coercion, a precedent showing how voice-synthesis technologies become weapons of political intimidation [3]. In Moldova and Taiwan, manipulated videos showed President Maia Sandu and candidate Lai Ching-te making statements that never existed, sparking protests and calls for recall that demanded a rapid institutional response.

These cases illustrate more than a simple "fact-check failure." The true danger is what experts call the **"liar's dividend"**: when the capacity to forge becomes ubiquitous, the public begins to doubt even verified evidence. This corrosion of faith in evidence creates a vacuum in which any claim — false or true — can be dismissed as potential manipulation. In this context, the defense of human rights faces a crisis of

credibility: international judges, criminal courts and investigative bodies find themselves unable to establish the authenticity of crucial evidence without advanced technological resources.

The solution does not lie only in detecting deepfakes after the fact, but in **authenticating communications from their origin**. Cryptographic hashing and immutable insertion on a blockchain provide a verifiable trail that any interested party can consult. When a military message or an official statement is registered with a unique hash at the moment of its creation, any later alteration is evident and automatically rejected by the system. Malicious actors are thus discouraged from creating forgeries, since documentary truth is "sealed" with a proof of existence that cannot be suppressed or manipulated.

In the end, the fifth front — defense and human rights — becomes a battle to preserve trust in information. Only by adopting protocols that guarantee the origin and integrity of data can we counter the liar's dividend and turn truth into a real weapon against disinformation.

False orders, impersonated commanders

In March 2024 a video circulated that appeared to come directly from the head of Ukrainian military intelligence, Kyrylo Budanov. Assembled from previous interviews with Budanov and with the secretary of the National Security and Defense Council, the clip fabricated a false "admission" attributing to Ukraine the attack on the Crocus City Hall concert venue. The material combined fragments of real voice and image with AI-generated manipulations until the result was almost indistinguishable from reality. The episode illustrates the imminent danger deepfake technology poses when a cloned voice or face is placed over a fabricated order or confession [4].

This vulnerability is not merely a problem of momentary confusion; it has strategic implications that can alter the balance of a conflict. A forged order can lead to unnecessary movements, expose forces to ambush and, in the worst scenario, trigger an accidental escalation of violence. That is why the joint guidance issued by the NSA, the FBI and CISA in 2023 identifies the Department of Defense and the defense industrial base as priority targets of impersonation attacks, and recommends two complementary axes: (1) detecting deepfakes with advanced forensic analysis tools, and (2) establishing the verifiable provenance of the media used to transmit orders and reports [4].

Establishing provenance relies on cryptographic identifiers ensuring that every message or document has been digitally signed by its legitimate issuer and that the resulting hash is anchored to an immutable blockchain. This model has already been put into practice in projects like **eyeWitness to Atrocities**, where an app generates a cryptographic hash in real time at the moment evidence is captured, guaranteeing its integrity and date of origin; more than 20,000 verifiable items have already been used as evidence in court cases [6]. Similarly, the **Starling Lab**, together with Hala Systems, has submitted evidence of war crimes to the ICC Prosecutor's Office using hashes, blockchain, C2PA and Filecoin, creating an auditable record that prevents any subsequent manipulation [7].

The risk extends beyond the terrestrial battlefield. The 2021 research on "deepfake geography" demonstrated how realistic-looking synthetic satellite images can be created, capable of faking the presence or absence of military activity in critical zones [5]. The director of the National Geospatial-Intelligence Agency (NGA) had already warned in 2019 that such deceptions could be used to conceal troop movements, manipulate diplomatic negotiations or justify armed interventions based on false evidence. If a government relies on a manipulated satellite image to plan an operation, the result can be a mistaken military response and the loss of innocent lives.

The solution, as the American authorities propose, is to integrate end-to-end authentication systems that combine digital signatures, blockchain anchoring and evaluation protocols run by independently certified agencies. Only with a verifiable chain of trust from the moment of generation to final consumption can we guarantee that military orders are not instruments of deception, but reliable instruments of a responsible defense. In this scenario, truth stops being a luxury and becomes the most powerful weapon against well-armed disinformation.

Truth as evidence: human rights

On an October afternoon in 2023, at a forward command center near Kharkiv, a team from the Starling Lab and Hala Systems received an apparently ordinary digital payload. The file included thousands of metadata items — timestamps, precise GPS coordinates, a cryptographic hash generated on the spot and a

link to the decentralized Filecoin network where the content would be stored. It was not just a document; it was the materialization of the **t2p** model, an immutable record from the very first instant, capable of blocking any attempt at manipulation.

The protocol built on the precedent set by the *eyeWitness to Atrocities* app, created in 2015 by the International Bar Association together with LexisNexis. Every image or video captured is instantly transformed into a cryptographic hash, accompanied by time metadata and GPS coordinates; that hash ascends into a verifiable chain of trust, and any alteration breaks the hash-file association, exposing the forgery with a clarity traditional courts cannot match [6]. To date, more than 20,000 items have been submitted to the United Nations Commission of Inquiry, becoming verifiable evidence in judicial proceedings related to the Ukrainian conflict.

When these elements were combined with the Starling Lab's blockchain infrastructure, the impact was exponential. The dossier presented to the Prosecutor of the International Criminal Court (ICC) in 2022 included evidence of alleged war crimes in Kharkiv: posts and videos of indiscriminate attacks, with the metadata registered and a hash for each item. Every piece was anchored in an immutable chain, preserved on Filecoin under the **C2PA** provenance standards and aligned with the **Berkeley Protocol**, adopted in 2020 as the global benchmark for digital evidence by bodies such as the OHCHR [7].

This architecture does more than guarantee authenticity; it creates a legal shield that blocks informational manipulation. At the ICC, a 2017 arrest warrant had already shown that evidence gathered from social networks could form the basis of a formal indictment; now, with hashes and blockchain, evidence reaches a level of certainty that overrides any doubt about its integrity. Truth ceases to be a subjective narrative and becomes **a legal weapon**: every fragment of reality is sealed from the moment of capture, unalterable all the way to the courtroom.

The message is clear: the *t2p* model is not merely a technical solution but a defense mechanism for human rights in the age of deepfakes and information warfare. Encoding truth with a hash and anchoring it to an immutable chain constitutes an impassable barrier against the forgery that could unleash panic or justify atrocities — and it has already proved able to safeguard justice and guarantee an authentic, unalterable voice for the victims.

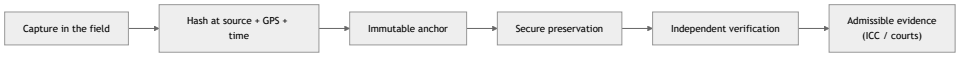


Figure 20.1 — Capture hashed at the origin → metadata and geolocation → immutable anchoring → admissible evidence: the chain that turns digital testimony into proof.

The defense responds

The intelligence agencies and NATO allies realized early that merely detecting deepfakes is only a first line of defense. The MediFor program, driven by DARPA, established a framework for validating the integrity of images and videos through advanced forensic analysis; but its success ran up against limits when generative neural networks began producing content indistinguishable from the original within days. In parallel, SemaFor — its semantic-forensics complement — culminated in 2024 with an open catalog of tools for the detection, attribution and characterization of manipulated material [8], and the NATO StratCom COE keeps continuous watch over new techniques of digital deception [8].

What these initiatives demonstrate is a fundamental principle: **after-the-fact detection cannot reverse the damage caused by falsified information that has already circulated**. As we saw with the Zelensky "surrender" deepfake of March 16, 2022, a fake video generated panic and influenced military operations before any forensic tool could confirm it [1]; similarly, the Russian "Doppelganger" campaign, which clones Western media websites with bots and AI to spread manipulated news, has been sanctioned by the EU yet continues to prove that **spread-before-identification** is the true engine of the harm [2].

That is why communications and evidence must be **authenticated from the moment of their creation**, a practice already proven in the human rights field. Applying the cryptographic hash at the origin, as the eyeWitness to Atrocities app does — generating a unique, immutable digest in real time during the capture of visual evidence — has made more than 20,000 items verifiable in court cases, setting a direct precedent for the "t2p" concept [6]. The model has scaled into initiatives like Starling Lab and Hala Systems, which presented documents of atrocities to the ICC Prosecutor with hashes, blockchain and the C2PA framework, turning the evidence into an unalterable, legally recognizable record [7].

The institutional response must evolve toward a **strategy of preventive sealing**: every military transmission, every intelligence report and every eyewitness account is encoded with a digital signature registered immediately on an immutable chain. That forecloses subsequent manipulation and gives judges and investigators a verifiable trail of provenance — evidence that cannot be talked away by the latest generative model.

In this context, the public challenge set by SemaFor in 2024 becomes a call to action: **develop and deploy open origin-sealing standards** compatible with the world's blockchain infrastructures. These protocols must be integrated into the command and control (C2) chain, ensuring that every order and every operational report carries a layer of cryptographic authenticity from the first byte.

Thus the most consequential front — defense against forgery in wartime — rests not merely on detecting lies, but on preventing lies from **prospering**. Truth, encoded and anchored to an immutable chain, becomes the most powerful weapon for preserving human rights and guaranteeing accountability in a world where information can be fabricated at the press of a button.

The promise and the peril

In conflict zones truth becomes a scarce good, but technology can turn it into an unstoppable weapon: cryptographic authentication guarantees that orders at the front are genuine and that evidence of atrocities can be presented in court without doubt as to its integrity. The *eyeWitness to Atrocities* project (IBA + LexisNexis) already generates a cryptographic hash at the moment of capture and links the metadata to an immutable chain; more than 20,000 verified items have been used as evidence in Ukraine, setting a precedent for the *t2p* framework [6]. Similarly, Starling Lab and Hala Systems have submitted evidence of alleged war crimes to the ICC Prosecutor with hashes and blockchain, using standards like C2PA and Filecoin; the Berkeley Protocol adopted by the OHCHR (2020) already envisaged this chain of trust [7].

This infrastructure not only strengthens the courts; it prevents military powers from constructing falsified pretexts to start or prolong conflicts. The Zelensky "surrender" deepfake, released on March 16, 2022 and debunked within hours, shows how an audiovisual manipulation can generate panic before the authorities react [1]; if every transmission were digitally signed, the impact of such forgeries would fall drastically. The same holds for the impersonation of commanders, like the deepfake of the Ukrainian intelligence chief Budanov (2024); the NSA/FBI/CISA recommendations (2023) for establishing provenance and authenticity in defense communications are now more relevant than ever [4].

Yet the threat bound to this capacity for authentication runs deep: encoding and anchoring metadata to an immutable chain can de-anonymize witnesses, victims and journalists in war zones if the proper safeguards are not in place. A record revealing the geographic location or exact time of a recording can cost lives, especially where authoritarian regimes have the resources to track and persecute dissidents. Such governments could also instrumentalize the technology to confer an "official seal" on fabricated propaganda, as the Russian "Doppelganger" campaign demonstrates — cloning Western media websites with bots and AI, sanctioned by the EU yet still capable of using verification technologies to legitimize falsehoods [2].

Balancing this duality demands a design centered on security and privacy. Selective disclosure should allow only authorized parties to access sensitive metadata, while preserving the transparency needed for judicial verification. The protocols must include witness-control mechanisms, such as temporary keys or decentralized credentials revocable in case of compromise. Only with this ethically conscious, mature framework can we harness the promise of an immutable truth without falling into the danger of totalitarian control of information.

It's time to act

The armed forces and security agencies can no longer rely on the voice of command alone; they must **authenticate every transmission** before it reaches the battlefield. The case of the deepfake of Ukrainian intelligence chief Budanov, which in 2024 tried to fabricate a false attribution of an attack, showed that the absence of a verification mechanism can trigger mistaken strategic decisions and raise the risk of escalation [4]. Forces must therefore incorporate digital signatures based on temporary keys and decentralized credentials, as the NSA, the FBI and CISA recommend for the provenance of military

communications (2023). Every order and every intelligence report must arrive with a **verifiable cryptographic hash** guaranteeing the message's origin and integrity before it is executed.

Human rights organizations and journalists are the front most vulnerable to manipulation. The *eyeWitness to Atrocities* app, developed by the IBA and LexisNexis, has proved that a **hash generated at the moment of capture** can provide an immutable proof of existence; more than 20,000 items have already served as reliable evidence in Ukrainian courts [6]. This model must be adopted globally: every photo, video or testimony recorded in conflict zones must be immediately sealed on a public, verifiable blockchain. Organizations must invest in mobile equipment integrating C2PA and Filecoin technology, as Starling Lab and Hala Systems did for the ICC in 2022 [7], ensuring the source is protected while the evidence cannot be forged.

Courts and prosecutors must go beyond passive acceptance. They must establish **mandatory protocols** based on the Berkeley Protocol (OHCHR, 2020) requiring hashed, blockchain-anchored evidence for any documentation related to war crimes. That means operational guides for judges and prosecutors, specialized training and a digital judicial infrastructure capable of validating evidence in real time. The first arrest warrant based principally on social media (al-Werfalli, 2017) already showed how decisive digital evidence can be; what is needed now is a formalized framework to guarantee its admissibility and authenticity.

The distribution platforms — from X to TikTok — must take proactive responsibility for preserving the provenance of conflict-related content; systems must be integrated that record metadata, timestamps and hashes at upload, with tools for independent verifiers. European regulation has already sanctioned impersonation campaigns like the Russian "Doppelganger" [2]; the platforms must go beyond simple blocking and guarantee full transparency about the material's origin.

Every user, as the first filter, must learn to verify before sharing images or audio from a war, using applications that display the hashes and the associated blockchain. The Zelensky surrender deepfake of 2022 — debunked within hours yet capable of feeding massive rumors — demonstrates how the absence of verification can create panic and legitimize violent action [1].

Policymakers must legislate clear standards for the security and safeguarding of these technologies, establishing ethical frameworks that prevent totalitarian control of information while fostering an immutable truth. **Time2act** is not just a motto: it is an urgent call for the proof of existence to become the foundation of justice and peace.

---

[1] NPR — *Zelenskyy deepfake telling Ukrainians to surrender* (16-Mar-2022). https://www.npr.org/2022/03/16/1087062648/deepfake-video-zelenskyy-experts-war-manipulation-ukraine-russia ↩

[2] EU DisinfoLab — *Doppelganger* (Russian disinformation campaign). https://www.disinfo.eu/doppelganger-hub/ ↩

[3] FCC — *Proposed $6M fine for AI-deepfake Biden robocall* (2024); IPI (Slovakia). https://www.fcc.gov/document/fcc-proposes-6m-fine-nh-robocalls ↩

[4] CISA / NSA / FBI — *Contextualizing Deepfake Threats to Organizations* (2023). https://www.cisa.gov/news-events/alerts/2023/09/12/nsa-fbi-and-cisa-release-cybersecurity-information-sheet-deepfake-threats ↩

[5] University of Washington — *Deepfake geography: AI falsifies satellite images* (Bo Zhao, 2021). https://www.washington.edu/news/2021/04/21/a-growing-problem-of-deepfake-geography-how-ai-falsifies-satellite-images/ ↩

[6] International Bar Association — *eyeWitness to Atrocities* (hash at capture; cases in Ukraine). https://www.ibanet.org/IBA-founded-eyeWitness-to-Atrocities-marks-ten-years-of-capturing-evidence-of-human-rights-crimes ↩

[7] USC Shoah Foundation — *Starling Lab & Hala Systems file cryptographic war-crimes evidence to the ICC* (2022); Berkeley Protocol (OHCHR). https://sfi.usc.edu/news/2022/06/33571-starling-lab-and-hala-systems-file-cryptographic-submission-evidence-war-crimes ↩

[8] DARPA — *Semantic Forensics (SemaFor)* and *Deepfake Defense* (2024); NATO StratCom COE. https://www.darpa.mil/research/programs/semantic-forensics ↩

Carbon Credits: The Compliance Market

"A core goal of public policy should be to facilitate the development of institutions that bring out the best in humans."

— Elinor Ostrom, Nobel lecture (2009)

On the morning of January 19, 2011, traders across half a continent found the same answer on their screens: *transfer unavailable*. It was not a computer failure. It was a deliberate decision by the European Commission, which that day pressed the off switch on the world's largest carbon market. Over the preceding weeks, attackers had emptied accounts at several national registries — Austria, the Czech Republic, Greece, Estonia and Poland — using *phishing* emails and password-stealing trojans. In a few days more than two million emission allowances disappeared, valued at some €28-30 million; the Czech registry was hit hardest, with losses of about €7 million [11]. Those allowances represented real tonnes of CO_2 that someone was legally obliged to offset, and they had flown off with a single click into accounts in Germany, Estonia and the United Kingdom.

The Commission had no choice but to suspend transfers at **all** the national registries at once and reopen them in stages, only as each met certain minimum security standards [11]. For days, Europe's principal climate instrument — a market that moves tens of billions of euros — sat frozen because its truth lived in databases that could be emptied in the small hours. The paradox was cruel: a system designed to save the climate had been hijacked by the most ordinary of thefts.

That shutdown is the best doorway into an uncomfortable question. If carbon has a price — and it does — who guarantees that the tonne being bought, sold or offset really existed, exactly once, and that no one has altered it since? This chapter is about the **compliance market**: the mandatory systems that put a price on emissions. They work, they move capital and they have already driven emissions down. But their integrity depends not on the price, but on something humbler and more fragile: trust in the data.

Putting a price on air

The heart of the European market is the emissions trading scheme (EU ETS), a cap-and-trade system. The idea is simple and elegant: the state sets a declining ceiling on the total gases that can be emitted and distributes or auctions allowances — the so-called EUAs, European Union Allowances — each equivalent to one tonne of CO_2 equivalent. Whoever pollutes less than their share can sell the surplus allowances; whoever pollutes more must buy them. The price rewards those who decarbonize and penalizes those who lag.

Twenty years after Paul Roberts, journalist and energy analyst, described the energy markets in *The End of Oil* (2004) as a great exchange where prices arbitrated between supply and demand, the ETS is perhaps the most ambitious market ever built on an abstraction: the right not to emit. As we have seen, it covers roughly 40% of the EU's greenhouse gas emissions and has raised more than €175 billion since 2013. The system — electricity, heat, heavy industry, aviation and, since January 2024, maritime transport — has helped cut the covered sectors' emissions by about 47% relative to 2005. The revised cap aims for a 62% reduction by 2030 against that year [1].

The price per tonne, for years too low to move anything, now bites; in mid-2026 it traded around **€80**, after nearly reaching **€106 in February 2023** [2]. To keep this signal from sinking when allowances are in surplus, the EU maintains a **Market Stability Reserve** that withdraws the excess from circulation [5]; and to prevent carbon leakage toward countries without a carbon price, 2026 brought onto the stage the newest and most debated piece: the **Carbon Border Adjustment Mechanism** (CBAM).

The CBAM entered its **definitive phase on January 1, 2026** [3], making imports of carbon-intensive products — iron and steel, aluminum, cement, fertilizers, electricity and hydrogen — pay the same CO_2 price that European industry already pays. The first quarterly CBAM certificate price, published on April 7, 2026, was set at **€75.36 per tonne**, calculated from the average of the first quarter's EUA auctions [4]. Certificate sales and surrenders do not begin until February 2027 and the first declaration falls due on September 30 of that year; a minimum threshold of 50 tonnes per year exempts 90% of small importers while keeping 99% of embedded emissions covered [4]. In parallel, the free allocation of allowances to industry will be phased out as the CBAM gains strength, and in aviation it already disappears in 2026 [5].

The **ETS2**, the new market intended to price the carbon of buildings and road transport, has been **postponed to 2028**: the political agreement was sealed in November 2025, within the framework of Europe's 2040 climate target, and formalization came in early 2026 [6]. Beyond Europe, international aviation has its own scheme, the **CORSIA** of the International Civil Aviation Organization: in 2024 alone, international flights emitted close to **597 million tonnes** of CO_2, and the sector was required to offset some **58 million** [7]. Three markets, three different designs; one and the same Achilles' heel.

The crime that hijacked the climate

For a market like this to exist, a tonne of CO_2 must be able to change hands as fast as a share on a stock exchange. Here lies the vulnerability: an emission allowance is a purely digital asset, with no weight or texture, transferred in seconds between accounts across Europe. That liquidity, virtuous for the market, proved an irresistible temptation for organized crime.

Between 2008 and 2010, a wave of **VAT carousel fraud** — what investigators call *missing trader fraud* — infected the market. The scheme was as old as commerce itself: a fraudster buys allowances in a country without VAT, sells them in another while charging the tax, and vanishes without remitting it to the treasury, repeating the operation in a circle. The European police agency **Europol** calculated that this fraud cost European taxpayers **more than €5 billion** and, as we have seen, in some countries the fraud could account for nearly the whole volume of the market [9]. The epicenter was the French exchange BlueNext, which offered immediate VAT refunds and saw activity there explode at the end of 2008. In Germany, prosecutors put the hole at some **€850 million** and, in December 2011, a court sentenced six people to between three years and seven years and ten months in prison [11].

The person who has best reconstructed that episode is **Marius-Christian Frunza**, a doctor of mathematics from the Sorbonne and former head of structuring at the commodities arm of France's Caisse des Dépôts, today a researcher of financial fraud. In his book *Fraud and Carbon Markets: The Carbon Connection* (Routledge, 2013) and the academic article *An Assessment of Carousel VAT Fraud in the European Carbon Market* (2014), Frunza demonstrates with forensic statistics that the **transfer speed of an intangible, untraceable asset** made it the perfect vehicle for the carousel [10]. The lesson is uncomfortable: a market built on simple bookkeeping entries, without a layer that seals the evidence, is vulnerable from the origin. The 2011 theft that opened this chapter was no anomaly; it was the second act of the same play.

The regulator and the chain of truth

Faced with these holes, the regulatory response has been to build an ever more demanding control architecture. The entire credibility of the ETS rests on one acronym: **MRV**, for *Monitoring, Reporting and Verification*. Two regulations sustain it, the Monitoring and Reporting Regulation (MRR) and the Accreditation and Verification Regulation (AVR), which oblige each installation to measure its emissions, draw up a report and have it validated by an independent **accredited verifier** before March 31 of each year [8].

Seen up close, all this machinery is, at bottom, a long **chain of data claims**. How many tonnes did that power plant emit? How much CO_2 is embedded in the steel crossing the border with a CBAM certificate? How much carbon did that flight burn? Each link in the chain is a document — a sensor reading, a spreadsheet, a PDF report, a certificate — that travels from hand to hand until it reaches the regulator. And each document is, at the same time, a point where the truth can break: through error, through negligence or, as we have seen, through malice.

The problem is that the traditional tools for protecting this chain — metadata, corporate seals, the annual audit, faith in the central registry — are precisely the ones digitalization has weakened. A PDF can be edited; a spreadsheet, retouched; a central registry can be hacked and then has to be switched off entirely. The integrity of the carbon market faces the same dilemma running through this whole book: in a world where evidence can be fabricated, **proving that a piece of data existed, intact, on a certain date** is no longer a formality — it is the foundation of everything.

Evidence that can be edited

Imagine an industrial installation preparing its annual emissions report. The technician downloads the sensor readings, pours them into a spreadsheet, applies the conversion factors and generates a PDF for the verifier to sign. Between the sensor reading and the verifier's signature lie a dozen steps at which the number can be "adjusted": a decimal here, an operating hour forgotten there. No conspiracy is required; the pressure to close the year with fewer tonnes than the real figure is enough, because every tonne saved is an allowance that can be sold.

The verifier does the job rigorously, but works with what is provided: documents that have already passed through interested hands and that nothing prevents from being touched up after the fact. If a dispute arises tomorrow — an inspection, a lawsuit, a suspicion of double counting between two national registries — how do you prove what the **original** sensor reading was, before any adjustment? The honest answer, today, is often that you cannot. The chain of custody of climate data is as weak as that of any digital file. And a market moving billions on data that cannot be audited back to the origin is a market that invites fraud.

The proof of existence as carbon's trust layer

This is where the **proof of existence** — the technology that runs through *time2act* via **t2p** — enters the scene, not as a silver bullet, but as a missing link. The idea is to attack the problem where it is born: instead of trusting that no one will touch the data afterward, the data is sealed at the exact moment it is generated.

The mechanism is the same one we know from other chapters, now applied to carbon. The verified MRV report, the raw reading of an emissions sensor, the calculation of a CBAM product's embedded emissions or the certificate of a CORSIA unit receives, at the origin, a **SHA-512 cryptographic fingerprint** computed on the device itself. That fingerprint is anchored to the public chain; as a complement, an **RFC 3161 timestamp** conforming to Europe's eIDAS rules, independent of the chain, adds the certification of an authority; from that instant, any later alteration of the PDF, the spreadsheet or the sensor log is detected immediately. For field audits — the inspection of an installation or the check of an offset project — **forensic GPS** and **proof of life** add another layer: they seal *where* and *when* the sample or the photograph was taken and tie the evidence to a real person present, making the classic desk-fabricated "office report" much harder.

And there is one more decisive detail, connecting directly to January 2011 — as we have seen. The evidence can be packaged in an **.asice container**, a format verifiable *offline*: a regulator, or a steel buyer subject to the CBAM, can check the integrity of the whole chain **without depending on the central registry**. That is, without the single point of failure that had to be switched off that day. Trust stops living in one vulnerable database and becomes checkable by anyone.

This is no marginal idea. The **World Bank**, through its Climate Warehouse program and the **Climate Action Data Trust** platform — developed with the government of Singapore and the international emissions trading association — is building exactly this: a common, tamper-proof data layer, supported by distributed ledger technology, that harmonizes information across national registries to prevent the **double counting** of a single credit [12]. The World Bank's thesis is crisp: the integrity of the post-Paris carbon markets depends on a shared data infrastructure that mere spreadsheet accounting can never provide [12].

Honesty about the limits is due, however, faithful to the duality of **"the promise and the peril"** that Don Tapscott, co-author of *Blockchain Revolution* (2016), uses to think about technology. The proof of existence guarantees that a piece of data has not been touched *after* sealing and, combined with a distributed ledger, helps prevent double counting. But it does **not** validate that the sensor was properly calibrated or that the initial measurement was true: "garbage in, garbage on-chain." Nor would it have stopped the VAT carousel fraud, which was a tax problem, not a data-integrity problem. The seal does not replace the audit or common sense; it only ensures that, when needed, the truth of the origin can be proven.

The promise and the peril

The promise is tangible. A carbon market with evidence sealed at the origin would be far more resistant to fraud, would make audits faster and cheaper, and would let a buyer, an investor or a citizen check for themselves that an offset tonne truly existed. In a system already moving more than €175 billion, as we

have seen, recovering this trust is no luxury: it is the condition for the carbon price to do the job it was invented for.

The peril, as always, calls for nuance. A badly designed technical layer can breed a false sense of security — the seal certifies the *when* and the *what*, not the physical truth of the measurement. Concentrating verification in few hands, or on a single platform, would recreate the very single point of failure we want to avoid. And total traceability, badly governed, can collide with companies' legitimate commercial confidentiality. The proof of existence is a **necessary** layer within a broader trust architecture, not a shortcut that spares us governance.

It's time to act

For operators of installations subject to the ETS: seal sensor readings and MRV reports at the origin, before any processing, so the original data is always provable. For importers subject to the CBAM: anchor the calculation of your products' embedded emissions, so that the figure declared at the border is auditable back to the factory. For accredited verifiers: demand proofs of existence as part of the chain of custody and refuse documents that could have been edited after generation. For registry managers and regulators: move toward a distributed, interoperable data layer — along the lines of the Climate Action Data Trust — that makes it impossible to switch off the whole market with a single click, and give the seal's evidentiary value legal recognition.

As we saw on 19/01/2011, Europe had to switch off its carbon market because the truth lived in places that could be emptied in the small hours. Fifteen years later we have the tools to prevent it; the carbon price will only work if we can trust every tonne that enters it, and proving that trust is time2act.

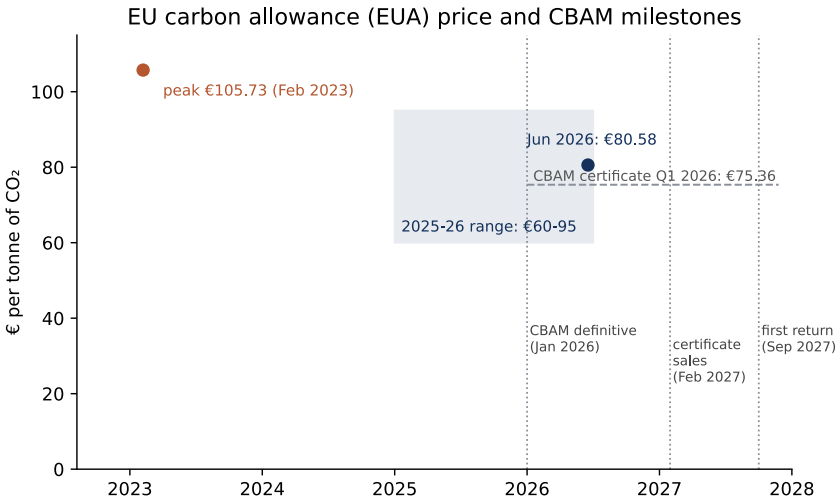


Figure 21.1 — The price of the European emission allowance (EUA) climbed to a peak of €105.73 in February 2023 and moved around €80 in mid-2026; in parallel, the CBAM entered its definitive phase on January 1, 2026, with a first certificate set at €75.36/t. Sources: Trading Economics; European Commission; S&P Global.

[1] European Commission — *About the EU Emissions Trading System* (coverage ~40%, €175 billion raised, −47% vs 2005, −62% by 2030). https://climate.ec.europa.eu/eu-action/carbon-markets/about-eu-ets_en ↩

[2] Trading Economics — *EU Carbon Permits* (EUA €80.58/t on 19-Jun-2026; all-time high €105.73 in Feb-2023). https://tradingeconomics.com/commodity/carbon ↩

[3] European Commission (Taxation and Customs Union) — *Carbon Border Adjustment Mechanism (CBAM)* (definitive phase from 1-Jan-2026). https://taxation-customs.ec.europa.eu/carbon-border-adjustment-mechanism_en ↩

[4] S&P Global — *Brussels confirms first CBAM certificate price for Q1 2026 at €75.36/tCO2e* (7-Apr-2026). https://www.spglobal.com/energy/en/news-research/latest-news/energy-transition/040726-brussels-confirms-first-cbam-certificate-price-for-q1-2026-at-eur7536mtco2e ↩

[5] International Carbon Action Partnership (ICAP) — *EU Emissions Trading System (EU ETS)* (MSR, CBAM factor, free allocation). https://icapcarbonaction.com/en/ets/eu-emissions-trading-system-eu-ets ↩

[6] European Commission — *ETS2: buildings, road transport and additional sectors* (start postponed to 2028; agreement Nov-2025). https://climate.ec.europa.eu/eu-action/carbon-markets/ets2-buildings-road-transport-and-additional-sectors_en ↩

[7] EASA — *ICAO CORSIA* (2024 international emissions ≈ 597 $MtCO_2$; ≈ 58 $MtCO_2$ to offset). https://www.easa.europa.eu/en/domains/environment/icao-carbon-offsetting-and-reduction-scheme-international-aviation-corsia ↩

[8] European Commission — *Monitoring, reporting and verification of EU ETS emissions* (MRR and AVR; verified report before March 31). https://climate.ec.europa.eu/eu-action/carbon-markets/eu-emissions-trading-system-eu-ets/monitoring-reporting-and-verification_en ↩

[9] Europol — *Carbon Credit fraud causes more than 5 billion euros damage for European Taxpayer* (up to 90% of volume fraudulent in some countries). https://www.europol.europa.eu/media-press/newsroom/news/carbon-credit-fraud-causes-more-5-billion-euros-damage-for-european-taxpayer ↩

[10] Marius-Christian Frunza — *Fraud and Carbon Markets: The Carbon Connection* (Routledge, 2013) and *An Assessment of Carousel VAT Fraud in the European Carbon Market* (Review of Law & Economics, 2014). https://www.routledge.com/Fraud-and-Carbon-Markets-The-Carbon-Connection/Frunza/p/book/9781138928091 ↩

[11] European Commission — *Announcement: transitional measure for the EU ETS registry system* (19-Jan-2011) and OCCRP — *Germany: Six Executives Stand Trial for Emissions Fraud* (~€850 million; six convicted). https://climate.ec.europa.eu/news-your-voice/news/announcement-transitional-measure-eu-ets-registry-system-2011-01-19_en · https://www.occrp.org/en/news/germany-six-executives-stand-trial-for-emissions-fraud ↩

[12] World Bank — *Why data infrastructure is key for a transparent carbon market* and the *Climate Action Data Trust* platform. https://blogs.worldbank.org/en/climatechange/why-data-infrastructure-key-transparent-carbon-market · https://climateactiondata.org/ ↩

The Voluntary Carbon Market: Promise, Fraud and the Race for Integrity

"If, like truth, the lie had but one face, we would be on better terms... But the reverse of truth has a hundred thousand shapes and a field indefinite."

— Michel de Montaigne, Essays, I, 9 ("Of Liars," 1580)

On January 18, 2023, three newsrooms — *The Guardian* in London, *Die Zeit* in Hamburg and the investigative collective *SourceMaterial* — simultaneously published the result of nine months of work. They had crossed satellite studies with fieldwork and academic science to answer an apparently technical question: do the carbon credits the big brands buy in order to declare themselves "neutral" really work? The conclusion shook an entire sector. More than **90% of the rainforest credits** certified by Verra — the largest standard in the world — were probably **"phantom credits"**: they represented no real reduction in emissions [1].

The detail that hurt was the list of buyers. Companies like Disney, Shell, Gucci and easyJet had advertised "carbon neutral" products and flights leaning on those credits. According to the analysis, the deforestation threat the projects claimed to avoid was overestimated on average by around **400%**, and only a small fraction of the credits examined corresponded to real reductions [1]. Verra rebutted the investigation harshly, calling it "manifestly unreliable" and accusing the studies of using "synthetic controls" that compared apples with oranges [2]. But the damage was done: four months later, on May 22, 2023, David Antonioli, Verra's founding chief executive, announced his departure after almost fifteen years at the helm [3].

That earthquake sums up the paradox of the voluntary carbon market. On paper, it is one of the most beautiful ideas in the climate fight: allowing those who cannot cut their emissions to pay for someone else to cut or absorb them — planting forests, conserving rainforest, capturing methane. But when the proof that this has really happened comes down to reports that can be inflated, the dream turns into an alibi factory. This chapter is about how to recover trust in a market that needs it desperately.

How the voluntary market works — and how it breaks

Unlike the compliance market of the previous chapter, here no one forces anything: companies buy credits **voluntarily**, to meet "net zero" promises or to polish their image. One credit nominally equals one tonne of CO_2 avoided or removed. Five certification programs — ACR, ART, Climate Action Reserve, Gold Standard and, above all, **Verra** — concentrate close to **98%** of the market [8]. Verra alone has at times controlled about two-thirds.

It is a small, volatile market. After peaking at around **$2 billion in 2021**, it contracted two years running to **$723 million in 2023**, with a **56% drop in transaction volume** relative to 2022 and an average price of barely **$6.53 per tonne** [8]. This price fragility has a root cause: no one knows for sure what a credit is worth, because no one can easily verify what stands behind it.

And behind it there must be four things, all hard to prove. **Additionality**: that the reduction would not have happened anyway without the credit's money. **Baseline**: how much deforestation or emission has *really* been avoided relative to a hypothetical scenario. **Permanence**: that the forest is still standing in ten or thirty years, not burned or felled. And **no double counting**: that the same tonne is not sold twice, nor counted simultaneously by the buyer and by the country where it is generated. Each of these properties is "certified," today, with documents, photographs and reports that can be manipulated. This is exactly where the fraud crept in.

The phantom credits scandal

The 2023 investigation did not come out of nowhere; it confirmed a suspicion science had already raised. **Patrick Greenfield**, environmental reporter at *The Guardian* and one of the work's authors, summed up its thesis: most of the Verra rainforest credits used by the big brands did not represent real reductions, and the logging risk they claimed to avert was systematically inflated [1]. It was no isolated opinion. A peer-reviewed study led by **Alejandro Guizar-Coutiño**, a researcher at the University of Cambridge, had evaluated forty REDD+ projects in nine countries in 2022 and found that they did reduce deforestation —

but only by **47%** on average: real reductions, yes, but far below what the credits sold had promised [6]. Another study, published in *Science* in 2023 on twenty-six sites in six countries — with detailed baseline analysis at eighteen — went further still, holding that most projects had not significantly reduced deforestation, though this work drew a formal reply and the choice of its counterfactual scenarios is still debated [7].

The case that best illustrates the problem is **Kariba**, in northern Zimbabwe: one of the world's largest REDD+ projects, covering some 785,000 hectares, which had sold credits to giants like Volkswagen, Gucci, Nestlé and McKinsey. When Verra investigated it thoroughly — prompted, in part, by a *New Yorker* feature — the certifier's own audit concluded that **15,220,520 credits had been issued in excess out of a total of 26,822,953**, that is, a **57% over-issuance** [4]. The project developer, the Swiss company South Pole, broke its contract with the local partner on October 27, 2023 [5]. Millions of "junk" credits were already circulating through the market, converted into certificates of a neutrality that never physically existed.

The mechanism of the fraud, at bottom, is always the same: you do not need to falsify the forest — inflating the *story* about the forest is enough. An exaggerated baseline scenario, a reused photograph, a measurement taken from an office chair, a report dated after a fire. All of it paper and pixels that no one can check back to the origin.

The race for integrity

Confronted with this discredit, the sector has launched a race to win back trust. The centerpiece is the **Integrity Council for the Voluntary Carbon Market** (ICVCM), an independent body chaired by **Annette Nazareth**, a former commissioner of the US SEC. In March the **Core Carbon Principles** (CCP) were published, a kind of quality seal that a credit must meet on permanence, robust measurement, verification and additionality. As Nazareth explained, the CCP label exists so that buyers can distinguish the credits that do meet rigorous requirements [9]. In June 2024, the ICVCM approved the first seven methodologies with the CCP seal, enabling some 27 million high-integrity credits [9]. In parallel, the **VCMI** initiative launched a code of good practice for the demand side — what a company buying credits may and may not claim [10].

At the level of states, **Article 6 of the Paris Agreement** tries to bring order to international trading. At the COP29 summit in Baku, in November 2024, the standards that make the United Nations mechanism operational were adopted [13]. The key is the **corresponding adjustment**: every credit authorized for international use obliges the selling country to subtract that reduction from its own inventory, so the tonne is counted only once. It is, in climate language, the solution to double counting — the same problem that, in the digital world, is solved by anchoring a unique fingerprint in a ledger no one can rewrite.

The temptation of blockchain done badly

A warning is due here that this book cannot spare itself. When someone says "blockchain," they often imagine a magic wand that makes everything transparent; the history of tokenized carbon proves the opposite: **badly applied technology made the problem worse**.

As we have seen, in October 2021 the Toucan and KlimaDAO projects launched on the Polygon network the *Base Carbon Tonne*, a token that promised to equal one tonne of CO_2. Within months, more than 21.6 million credits migrated onto the chain. The analysis center **CarbonPlan** uncovered the problem: its researchers **Danny Cullenward and Grayson Badgley**, in the study *Zombies on the blockchain* (April 2022), showed that nearly 28% of those credits came from "zombie" projects — old, low-quality credits that careful buyers had already discarded, and that tokenization had resurrected [11]. The blockchain had become, they said, a dumping ground where the worst credits found a second life. The flagship token reached almost $3,900 and ended up sunk at around $20. In May 2022, Verra itself banned the tokenization of already retired credits [12].

The lesson is exactly the duality of **"the promise and the peril"** with which Don Tapscott, co-author of *Blockchain Revolution* (2016), insists on thinking about every technology. Putting something "on the chain" does not make it true; it only makes it *immutable*. If what you anchor is garbage, you will have immutable garbage. The right question is not *whether* we use the blockchain, but **what we anchor in it and at what moment**.

The proof of existence done right: sealing evidence, not inventing it

It is on this question that the **proof of existence** of **t2p** marks the difference from the failed experiments. The idea is not to tokenize the final credit — the product already contaminated by a dubious methodology — but to **seal the primary evidence** at the very instant and place it is generated, before anyone can retouch it.

In practice: every photograph of a reforestation, every measurement of an inventory plot, every reading from a field sensor receives at the origin a **SHA-512 fingerprint** computed on the device, anchored to the **public blockchain** with an **RFC 3161 timestamp** and accompanied by **forensic GPS** and the technician's **proof of life**. It is thus demonstrated that *that* photograph was taken *at those coordinates*, *on that date*, by *a real person present*: not lifted from Google Earth, not generated by AI, not backdated after a fire. The entire evidence package of a monitoring round — photos, measurements, coordinates, signature — can be closed into an **offline-verifiable .asice container** that an auditor, a rating agency like Sylvera or BeZero, or the registry itself can check **without depending on whoever sells the project**, even years later. These fingerprints can be chained into open **digital MRV** systems, like Open Forest Protocol, and into the traceability the Article 6 corresponding adjustment demands [14].

One must be scrupulously honest about the limit, though, and it is an important one. The proof of existence **does not solve additionality and does not correct an inflated baseline** — which is the heart of the Verra and Kariba scandal. If the methodology is bad, the seal will only certify, with great cryptographic fidelity, that someone measured something in a place at a moment. "Garbage in, garbage on-chain." What t2p contributes is not validation of the counterfactual assumption, but **raising confidence in the field evidence**: it makes the reused photo, the desk measurement and the postdated report impossible. Climate integrity still depends on good baseline science and independent auditing. The technology is a necessary condition, not a sufficient one.

The promise and the peril

The promise is powerful. A voluntary market where every piece of evidence is sealed at the origin would be far harder to cheat, would let a buyer check on their own that the forest they finance exists and endures, and would restore credibility to credits that today sell at a discount because of distrust. In a sector that has lost more than half its volume in a single year, recovering credibility is, literally, a matter of survival.

The peril has three faces. First, the temptation to believe that technology replaces science: no seal compensates for a bad methodology. Second, the inverse placebo effect — a company buying "cryptographically verified" credits to avoid the hard work of cutting its own emissions; the seal certifies the evidence, but does not excuse decarbonizing. And third, concentration: if verification and rating end up in very few hands, we recreate the single point of trust we wanted to dissolve. The proof of existence is a tool in the service of integrity, not a substitute for integrity.

It's time to act

For project developers: seal the field evidence at the origin — photos, measurements, coordinates — so the project's quality is demonstrable and need not be taken on faith. For buying companies: demand credits with anchored evidence and an integrity label (CCP), without using them to postpone reducing your own emissions. For the registries and standards (Verra, Gold Standard, ICVCM): integrate the proof of existence and digital MRV into your methodologies, closing the door on the tokenization of zombie credits. For the rating agencies: make offline-verifiable evidence the basis of your grades. For regulators: make the Article 6 corresponding adjustment operational with tamper-proof data infrastructure.

The voluntary market, born of a generous idea, has crashed against its own opacity; its years are numbered unless it proves what it promises. As we have seen, the difference between a phantom credit and a real credit is not the rhetoric but the proof — and generating that proof before anyone can inflate it is **time2act**.

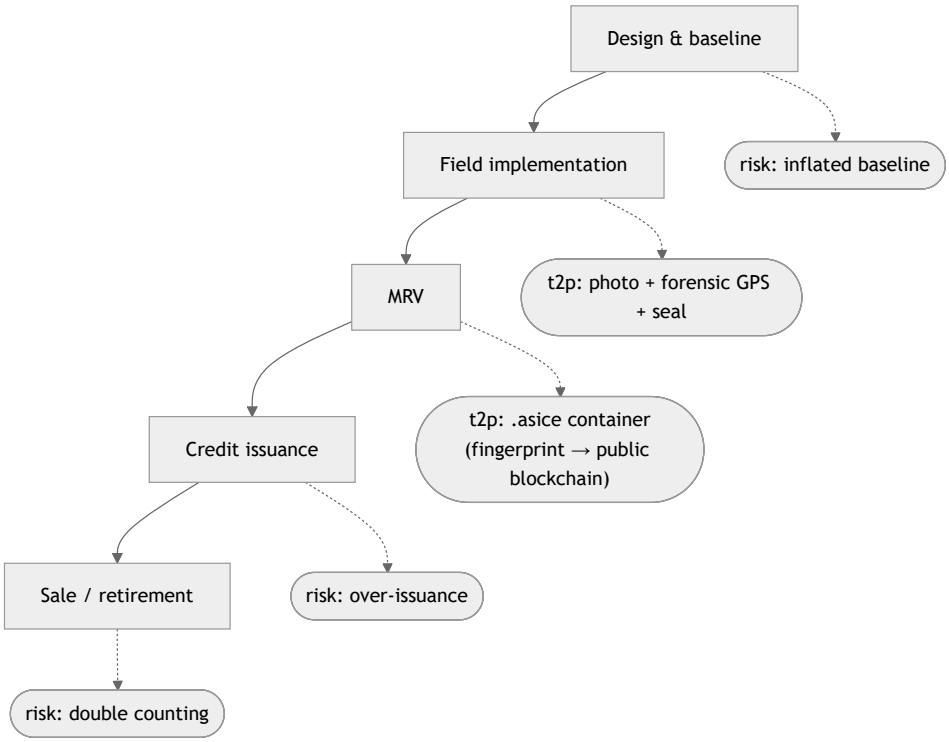


Figure 22.1 — Across a credit's life — design and baseline, implementation in the field, MRV, issuance, sale and optional tokenization — fraud infiltrates above all through the inflated baseline and the manipulable field evidence. The proof of existence anchors the evidence links (photo, measurement, report) but does not fix a defective methodology. Sources: The Guardian; Verra; CarbonPlan; ICVCM.

---

[1] The Guardian / Die Zeit / SourceMaterial — Patrick Greenfield, *Revealed: more than 90% of rainforest carbon offsets by biggest certifier are worthless* (18-Jan-2023). https://www.theguardian.com/environment/2023/jan/18/revealed-forest-carbon-offsets-biggest-provider-worthless-verra-aoe ↩

[2] Verra — *Verra Response to Guardian Article on Rainforest Offset Credits* (rebuttal). https://verra.org/verra-response-guardian-rainforest-carbon-offsets/ ↩

[3] Mongabay — *Head of Verra, top carbon credit certifier, to leave in June* (resignation of David Antonioli, 22-May-2023). https://news.mongabay.com/2023/05/head-of-verra-top-carbon-credit-certifier-to-leave-in-june/ ↩

[4] Verra — *Verra Acts on Kariba Project: Cancels Excess Credits and Advances Independent Review* (15,220,520 credits in excess out of 26,822,953; ~57%). https://verra.org/verra-acts-on-kariba-project-cancels-excess-credits-advances-independent-review/ ↩

[5] South Pole — *Statement on the Kariba REDD+ project* (27-Oct-2023). https://www.southpole.com/news/statement-27october ↩

[6] Alejandro Guizar-Coutiño et al. — *A global evaluation of the effectiveness of voluntary REDD+ projects* (Conservation Biology, 2022; deforestation −47%). https://pmc.ncbi.nlm.nih.gov/articles/PMC10086997/ ↩

[7] West et al. — *Action needed to make carbon offsets from forest conservation work for climate change mitigation* (Science, 2023; 26 sites in 6 countries; disputed result, arXiv reply 2312.06793). https://www.science.org/doi/10.1126/science.ade3535 ↩

[8] Ecosystem Marketplace / Forest Trends — *State of the Voluntary Carbon Market 2024* (market ~$2 billion in 2021 → $723 million in 2023; volume −56%; average price $6.53/t). https://www.ecosystemmarketplace.com/articles/report-the-voluntary-carbon-market-contracted-in-2023-driven-by-drop-off-in-transactions-for-redd-and-renewable-energy/ ↩

[9] ICVCM — *Core Carbon Principles* and *First high-integrity CCP-labelled carbon credits* (June 2024; ~27 million credits). https://icvcm.org/core-carbon-principles/ · https://icvcm.org/integrity-council-announces-first-high-integrity-ccp-labelled-carbon-credits-as-assessments-continue/ ↩

[10] VCMI — *Claims Code of Practice* (28-June-2023). https://vcmintegrity.org/launch-claims-code/ ↩

[11] CarbonPlan — Danny Cullenward and Grayson Badgley, *Zombies on the blockchain* (7-Apr-2022; ~28% zombie credits). https://carbonplan.org/research/toucan-crypto-offsets ↩

[12] Verra — *Verra Addresses Crypto Instruments and Tokens* (ban on tokenizing retired credits, 25-May-2022). https://verra.org/verra-addresses-crypto-instruments-and-tokens/ ↩

[13] S&P Global Commodity Insights — *COP29: rules for UN-led carbon market under Article 6.4 approved in Baku* (11-Nov-2024). https://www.spglobal.com/commodity-insights/en/news-research/latest-news/energy-transition/111124-cop29-rules-

for-un-led-carbon-market-under-article-64-approved-in-baku ↩

[14] CIFOR-ICRAF / Open Forest Protocol — *Blockchain platform offers reforestation efforts access to carbon finance* (open digital MRV). https://forestsnews.cifor.org/76905/blockchain-platform-offers-reforestation-efforts-access-to-carbon-finance ↩

CAE: Certifying Savings Without Cheating

"If it is not right, do not do it; if it is not true, do not say it."

— Marcus Aurelius, Meditations, XII, 17

In Pordenone, in the northeast of Italy, there was an energy services company that on paper was a model of efficiency. Since 2013 it had submitted to the Italian public manager, the GSE, nearly three hundred energy-saving projects: boilers replaced, insulation installed, equipment renewed in the homes of hundreds of customers. The GSE validated them and the company collected **50,322 white certificates**, the titles that reward every kilowatt-hour saved, with a profit of more than ten million euros.

There was just one problem. The company had no offices, no logistics, no employees. And the works did not exist. The "customers" were names lifted from the lists of a thermo-hydraulic supplies distributor and a solar panel dealer — people who had never ordered anything. The Guardia di Finanza uncovered it, and the regional Court of Auditors put the damage to the public purse at some **€10.4 million** [10]. It was the perfect fraud in its simplicity: certificates for savings measured with millimetric precision... on installations that existed only in a PDF.

The Pordenone case condenses the fragility of one of the most intelligent instruments of European energy policy: **energy savings certificates**. The idea is brilliant — paying for the energy that is *not* consumed, because the cleanest kilowatt-hour is the one that never needs generating. But the whole edifice rests on a humble proof: that the savings measure was really carried out, where it says and when it says. And that proof, until now, was a handful of papers and photographs that anyone could invent.

Paying for the energy we do not spend

It all starts from a European obligation. The **Energy Efficiency Directive** (2023/1791, Article 8) obliges each state to achieve new final energy savings year after year: 0.8% annually between 2021 and 2023, 1.3% in 2024-2025, 1.5% in 2026-2027 and up to **1.9% a year in 2028-2030**. The Commission itself acknowledges that the recast directive "more than doubles" the savings obligation toward 2028 [1]. For Spain, the decade's cumulative final energy savings target reaches **53,593 ktoe** [5]. It is a mountain of savings that someone has to generate and, above all, to prove.

The instrument that makes it possible is the **white certificate**, known by different names in each country. In Spain they are the **CAE**, the Certificados de Ahorro Energético (Energy Savings Certificates), created by **Royal Decree 36/2023** of January 24. The golden rule is as clear as a marked kilometer: **one CAE equals one kilowatt-hour of final energy saved and accredited** [2]. The system distributes the effort among well-defined actors: the **obligated parties** — gas and electricity retailers and oil product operators — must generate savings proportional to their sales; they can do it directly, through a **delegated party**, or by paying into the National Energy Efficiency Fund. Each measure must be validated by a **verifier accredited** by ENAC, and the whole system is coordinated by the Ministry for the Ecological Transition through a catalog of standardized measures [2][3].

The model works. In its first full year, 2024, the Spanish CAE system mobilized more than **1,200 measures**, certified close to **2 TWh** of annual savings, avoided some **500,000 tonnes of CO_2** and generated around **€230 million** in certificate sales, with industry as the main protagonist (68% of the savings) [4]. France has used the same principle since 2006 with its **CEE**, counted in a unit of their own, the "kWh cumac" — energy saved over the equipment's entire useful life — with an obligation that for the 2026-2030 period reaches **1,050 TWh cumac per year** [6]. Italy has its certificati bianchi. Three countries, three names, one and the same Achilles' heel.

The person who has studied these schemes best is **Louis-Gaëtan Giraudet**, an economist at CIRED and the École des Ponts and one of Europe's authorities on white certificates. In works like *Selling and Saving Energy: Energy Efficiency Obligations in Liberalized Energy Markets* (2020), Giraudet concludes that these systems usually show a positive cost-benefit balance, but that their integrity **depends critically on how savings are measured and verified** against the baseline [12]. In other words: the instrument is worth exactly what its proof is worth.

Fraud attacks the proof, not the formula

Here lies the crux. The energy-savings fraudster almost never disputes the kilowatt-hour formula; what he falsifies is **the proof that the physical measure exists**. Phantom installations, savings inflated against the baseline, fake invoices, double counting and non-additional measures that would have been carried out anyway. Every attack converges on the same weak point: the evidence.

France is the most mature and painful case study. The CEE scheme gave rise to the phenomenon of "isolation à 1 euro" — insulation for one euro: companies that recruited householders by phone promising to insulate their homes "for one euro," often claiming a mandate from official bodies. The result was poor materials, botched or outright nonexistent works and inflated certificate applications. In the summer of 2018 the French consumer and fraud authority, the DGCCRF, inspected 469 establishments in the sector and found that 56% were not fully complying with the rules [8]. The drip has continued: as we have seen, the French government has blocked €480 million of energy-transition fraud, of which €236 million corresponds to the CEE [9]; and in February 2025 the judicial customs service dismantled a network that had diverted thirteen million through fictitious works and transfers to tax havens.

Italy has answered with a revealing legal doctrine. Its highest administrative court, the Consiglio di Stato, has confirmed that the award of a white certificate is a "process of progressive formation": even if the GSE initially validated it, it retains the power-and-duty to check afterward that the savings really existed and, if the beneficiary cannot prove it with documentation, to revoke the certificate and demand its repayment [11]. In other words: an issued certificate **is not an untouchable right**. The burden of proof rests, forever, on whoever collected it; and if the proof is a folder of photos no one can authenticate, that burden is smoke.

The regulator is already demanding forensic proof

The most interesting development is that regulators have already arrived, by force of events, at the conclusion that runs through this book: trust must be provable. The French response to the fraud illustrates it almost literally. **Law 2025-594 of June 30, 2025** did not merely raise the ceiling on penalties — from 4% to **10%** of turnover, and up to 12% for repeat offenders; it also required certain savings operations to be justified with **dated and geolocated photographs** — *horodatées et géolocalisées* — or with remote video, evidence that must be kept for **six years** and made available to control agents [7].

The French state, after years of fake certificates, has stopped believing the installer's word and requires by law a **photograph with verifiable date and coordinates** for each measure. That is, precisely, a proof of existence. Third-party field checks rose from 10% to 15% between 2023 and 2025 [9]; the regulator has described the remedy, but the guarantee is still missing that the photograph, once taken, cannot be manipulated or trick-dated.

From legal requirement to unalterable trail

And this is where the **proof of existence** of **t2p** turns a good regulatory intention into a technical guarantee. French law asks for a dated, geolocated photo; the problem is that a date and coordinates, in an ordinary file, can also be faked. The proof of existence closes that crack by sealing the evidence at the exact moment and place it is generated.

The mechanism, applied to a savings measure — a boiler replacement, insulation, an efficient industrial motor, LED lighting — is as follows. The complete evidence is captured: *before* and *after* photos, initial and final meter readings, the installer's invoice, the equipment's serial number and the verifier's signature. Each file receives on the device a **SHA-512 fingerprint**, anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS and with **forensic GPS** fixing the coordinates of the work. The "after" photo also incorporates a **liveness check** to prevent the reuse of old or catalog images. Everything is packaged into an **.asice container verifiable offline**.

The result attacks every fraud vector. Geolocation and the liveness check force physical presence at the declared address: goodbye to Pordenone's phantom installations. The timestamp and the fingerprint expose any photo reused from another job or a catalog. Any later retouching of an invoice or a certificate breaks the anchored fingerprint, and the same fingerprint cannot be credited twice, which shuts the door on double counting. What French law already requires — a dated, geolocated photo kept for six years — t2p turns into an **unalterable audit trail**, verifiable by the Spanish Coordinator, by the French control

authority or by the Italian GSE **without having to trust the obligated party's own archive**, and at a verification cost far below field sampling.

As always, honesty about the limit is due. Anchoring a proof **does not certify that the savings are real**: if the accredited verifier endorses an inflated measure, or the baseline is miscalculated, the chain will only guarantee that that document existed on that date and in that place — not its technical truthfulness. "Garbage in, garbage on-chain." The proof of existence armors the **integrity of the evidence and the identity of whoever signs it**, shifting fraud from the "forgeable document" to the "expert who knowingly lies," a conduct far harder, more traceable and carrying personal liability. It does not replace the verifier or the physical measurement; it gives them a reliable foundation.

The promise and the peril

The promise is direct and quantifiable. A savings certificate system with evidence sealed at the origin would be far more fraud-resistant, would make verification cheaper, would speed up payment of legitimate certificates and would protect the billions of public and private euros these schemes move. In an instrument where every badly spent euro is a euro that buys no real savings, the integrity of the proof is not bureaucracy: it is climate effectiveness.

The peril demands nuance, in line with the duality of "the promise and the peril" that runs through this whole book. A technical layer can breed a false sense of security: the seal certifies the *when* and the *where*, not the quality of the measure. The geolocation and the biometrics of the liveness check touch sensitive data of installers and residents, and demand a privacy-respectful design. And if verification concentrates on a few platforms, we recreate the single point of trust we wanted to dissolve. The proof of existence is a tool in the service of certificate integrity, not a substitute for the verifier or for good baseline design.

It's time to act

For the **obligated and delegated parties** (retailers, energy service companies): seal the evidence of every measure at the origin, so the quality of the work is demonstrable and the certificates unassailable. For the **accredited verifiers**: demand proofs of existence as the basis of validation, and refuse photographs or invoices that could have been edited afterward. For the **coordinators and control bodies** (MITECO, the French PNCEE, the Italian GSE): recognize the seal's evidentiary value, integrate it into the requirements of the measures catalog and thereby cut the cost of sampling. For **European regulators**: harmonize the forensic-proof requirement France has already written into law, before every state reinvents the wheel.

In Pordenone, a company with no offices collected fifty-odd thousand certificates for works that did not exist. The technology to make that fraud impossible already exists, and an entire country is beginning to demand it by law. Certifying savings without cheating — proving that every kilowatt-hour saved truly existed — is, here too, **time2act**.

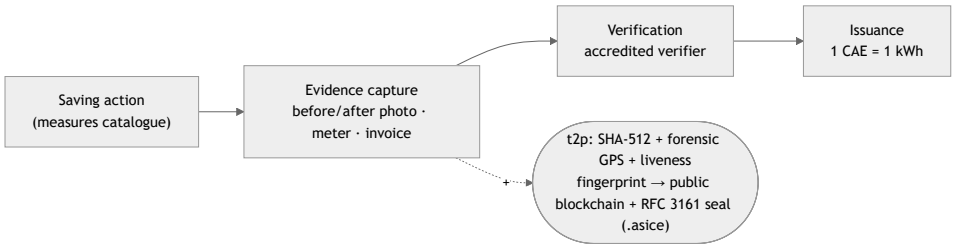


Figure 23.1 — From the measures catalog to certificate issuance: the measure is captured as evidence (before/after photo, meter, invoice), verified and converted into a CAE (1 CAE = 1 kWh). The proof of existence is inserted at the moment the evidence is captured — SHA-512, forensic GPS, RFC 3161 seal, liveness check and .asice container — exactly where French law already requires dated, geolocated photographs. Sources: BOE (RD 36/2023); Légifrance (Law 2025-594); Funcas.

[1] European Commission — *Energy Efficiency Directive* (Directive 2023/1791, Art. 8; new savings up to 1.9% annually in 2028-2030; "more than doubles" toward 2028). https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive_en ↩

[2] BOE — *Royal Decree 36/2023, of January 24* (CAE system; 1 CAE = 1 kWh of accredited final energy; obligated/delegated parties, verifier, coordinator). https://www.boe.es/buscar/doc.php?id=BOE-A-2023-2027 ↩

[3] BOE — *Order TED/845/2023, of July 18* (catalog of standardized energy efficiency measures). https://www.boe.es/buscar/doc.php?id=BOE-A-2023-16940 ↩

[4] Funcas — *Sistema de Certificados de Ahorro Energético (CAE): el motor que ya impulsa la transformación energética en España* (2024 results: ~1,200 measures, ~2 TWh, ~500,000 t CO_2, ~€230 million; industry 68%). https://www.funcas.es/articulos/sistema-de-certificados-de-ahorro-energetico-cae-el-motor-que-ya-impulsa-la-transformacion-energetica-en-espana/ ↩

[5] MITECO — *Memoria de la Orden de Obligaciones del Sistema CAE 2025* (cumulative savings target 2021-2030: 53,593 ktoe). https://www.miteco.gob.es/content/dam/miteco/es/energia/files-1/es-ES/Participacion/Documents/aeip-fnee-2025/Memoria%20OM_Obligaciones%202025_TA.pdf ↩

[6] Légifrance — *Décret n° 2025-1048 du 30 octobre 2025* (CEE, 6th period 2026-2030; 1,050 TWh cumac/year). https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000052486193 ↩

[7] Légifrance — *Loi n° 2025-594 du 30 juin 2025* (penalties up to 10-12%; photographs *horodatées et géolocalisées* or remote video, kept 6 years). https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000051824277 ↩

[8] DGCCRF (economie.gouv.fr) — *Rénovation énergétique : focus sur l'isolation à 1 euro* (469 establishments inspected; 56% non-compliance). https://www.economie.gouv.fr/dgccrf/laction-de-la-dgccrf/les-enquetes/renovation-energetique-focus-sur-lisolation-1-euro ↩

[9] Hellio / French government — *Fraude aux CEE* (€480 million of fraud blocked in 2024, €236 million in CEE; control rate 10%→15%). https://www.hellio.com/actualites/reglementation/fraude-cee ↩

[10] ANSA / PordenoneToday — *Contributi per 10 milioni su certificati bianchi inesistenti* (292 fictitious projects; 50,322 TEE; damage ~€10.4 million). https://www.ansa.it/sito/notizie/cronaca/2025/06/23/contributi-per-10-milioni-su-certificati-bianchi-inesistenti_dd57abfb-77ba-483a-8f42-c240874fc308.html ↩

[11] EnergiaMercato / Consiglio di Stato — *GSE and the annulment of energy efficiency titles* (the "progressive formation" doctrine; 4ESCO case). https://www.energiamercato.it/notizie/project-finance/gse-conferma-la-legittimita-dellannullamento-di-titoli-di-efficienza-energetica-a-4esco ↩

[12] Louis-Gaëtan Giraudet, Matthieu Glachant and Jean-Philippe Nicolaï — *Selling and Saving Energy: Energy Efficiency Obligations in Liberalized Energy Markets* (The Energy Journal, 2020). https://ideas.repec.org/f/pgi337.html ↩

Self-Consumption and Energy Communities: The Verifiable Reading

"The sun shines to-day also."

— Ralph Waldo Emerson, Nature (1836)

The afternoon the bill arrives, the line item reads: "surplus compensation — €0." On the other screen, the inverter's app shows a whole month of sunshine accounted for: a figure that looks nothing like the retailer's. The prosumer looks at the two windows, side by side: two realities that do not match, and no way of knowing which one rules.

This is no niche anecdote [1]: Spanish photovoltaic self-consumption has reached 9.3 gigawatts of installed capacity — the country's fifth-largest source of electrical generation by capacity, ahead of nuclear. In 2025 alone, 36,330 households joined, within a total of 1,139 new megawatts counting the commercial and industrial sectors too; and the national energy and climate plan's target is to reach 19 gigawatts [1]. Millions of families and small businesses have put their savings on the roof — and have become, without quite realizing it, generators with their own energy accounting.

This chapter closes a circle the book opened in Part II: we began with the energy machine, and we now meet it again on the rooftops — because the energy transition is also a transition of trust. Energy savings certificates — the measures that get paid — are in chapter 23; the residents' association minutes, in chapter 37; citizen environmental sensors, in chapter 97; the carbon markets, in chapters 21 and 22. The key question: who holds the meter when the bill does not match the sun you generated? The answer: the other party — unless your own reading has a certain date: the inverter's and the meter's series sealed in chained batches with **t2p** — the SHA-512 fingerprint computed on the device, the anchoring of the fingerprint to the public blockchain and the complementary RFC 3161/eIDAS timestamp, inside an .asice container [3]. And with the usual symmetry: the same trail disarms the unfounded claim against the retailer who settles correctly.

The other party's meter

The framework is clear and, in general, well designed [2]: Royal Decree 244/2019 regulates self-consumption and creates the simplified compensation of surpluses — voluntary, for renewable installations up to 100 kW: the energy fed into the grid is compensated financially on each month's bill (monthly net billing, not annual net metering). Self-consumption can be individual or collective — several consumers associated with a single installation, with distribution coefficients — and recent rules have extended its radius to 5 kilometers for plants of up to 5 megawatts [2]. On this foundation, energy communities are growing: neighbors, town councils and small businesses grouped together to share generation.

The problem is not the framework: it is the proof. The settlement of surpluses is made on the readings of the fiscal meter and the retailer's systems. The prosumer has the inverter — which measures production in real time — and its app; but when the two figures diverge, or the compensation takes months to appear, or a change of retailer loses the thread, the claim is argued against someone else's data: the other party's.

In the collective model, the uncertainty multiplies: the coefficients agreed among neighbors, the joiners and leavers, the assembly resolutions — often consigned to a goodwill spreadsheet. And in long-term power purchase agreements, the agreed prices and volumes must be checked against those settled — one more documentary layer.

The conclusion is this book's: the energy transition has decentralized generation, but not the proof. Whoever generates at home deserves energy accounting with a certain date of their own — not out of systematic distrust of anyone, but so that when the numbers do not add up, the conversation is one of facts.

The proof of existence applied

Four gestures — most of them automatic — turn the roof into a generator with verifiable accounting [3].

1. Your own reading, in chained batches

The periodic exports from the inverter and your own meter — production, instantaneous self-consumption, surpluses — are certified with t2p in batches: each closing, weekly or monthly, receives its SHA-512 fingerprint, chained to the previous one [3]. It is the same logic the state already imposes on invoices — chapter 34 explains it — applied to your own roof. The result: a dated series no one can rewrite — not even the prosumer, which is exactly why it counts before third parties.

2. The bill, cross-checked

Every bill, captured with a certain date on arrival; the compensation line, checked against your own series for the same period [3]. When they do not match, the claim carries dated numbers: these surpluses my inverter counted, this is what you compensated me, this is the difference. And the symmetry: the retailer who settles correctly has, in its records and in the client's series, the proof that its bill was right — the unfounded claim dies of dates.

3. The energy community's minutes

The distribution coefficients, each participant's joining and leaving, the assembly's resolutions — sealed with a certain date, like the residents' association minutes of chapter 37 [3]. When a participant disputes their percentage, or someone joins mid-year, the history of the agreements will be a chain of anchored versions, not an argument of memories. Energy communities will only scale if trust among strangers does not hang on the goodwill of a spreadsheet.

4. The frozen contract

The compensation contract, the retailer's offer on the day of signing, the PPA — the long-term power purchase agreement — with its prices and volumes: captured and sealed at the moment [3]. When the terms "change" or the fine print gets reinterpreted, that day's document exists, exact and dated. The same gesture chapter 33 applies to the phone tariff, applied to the longest contract a family signs.

The limits, stated plainly: the seal certifies series and documents — it does not settle, does not set tariffs and does not replace legal metrology: the fiscal meter is what it is, and your own reading contrasts, it does not substitute. A badly calibrated inverter seals its own deviation — maintenance and calibration first. And small discrepancies are normal — losses, measurement instants: the trail is for the large and the persistent ones.

The promise and the peril

The promise: prosumers who claim with their own dated series [1][3]; energy communities that grow because their agreements carry a certain date; diligent retailers who resolve claims in days because both parties hold verifiable numbers; and the country's fifth-largest generation source [1] with the trust accounting it deserves.

The perils, faced squarely.

The first is **the war of meters**. If each side brandishes its own series, we have gained nothing. The answer is methodological: calibrated equipment, declared methodology, and one's own reading as documented contrast that forces a review of the fiscal meter — not as a substitute.

The second is **prosumer overload**. The family that installed panels to save money does not want cryptographic bookkeeping. The sealing must be automatic — the periodic export that anchors itself [3] — or it will not happen; the manual gesture is for the day of the discrepancy, not for every day.

The third is **the energy profile as intimacy**. A household's consumption curve tells when you sleep, when you cook, when you are away. The series stay on the device or in the prosumer's systems; on the chain, only fingerprints [3]; disclosure, selective and to whom it may concern.

And the fourth is **solar techno-solutionism**. No hash speeds up a grid-access permit, expands a saturated network or fixes slow regulation. The proof of existence makes settlements fair; getting self-consumption to 19 gigawatts [1] is the task of energy policy.

It's time to act

Let every kilowatt generated at home have its certain date.

For prosumers: switch on the inverter's sealed periodic export [3]; the day the bill does not add up, you will have your series — and not just your word.

For energy communities: coefficients, memberships and resolutions with a certain date [3] — trust among neighbors deserves the same trail as any contract.

For retailers and distributors: anchor your settlements and publish them by fingerprint — the unfounded claim dies of dates, and your diligence becomes a commercial argument.

For the regulator and consumer authorities: recognize the anchored self-reading as qualified contrast in compensation disputes [2]; the prosumer cannot litigate blind against the other party's meter.

For legislators: the decentralized energy transition — already brushing 9.3 gigawatts and aspiring to 19 [1][2] — needs its trust layer: let the rollout of energy communities be born with verifiable minutes and settlements as standard.

This book began with the end of easy oil; it meets energy again now on the rooftops. Let the energy transition also be a transition of trust. For the sun on every roof, **time2act**.

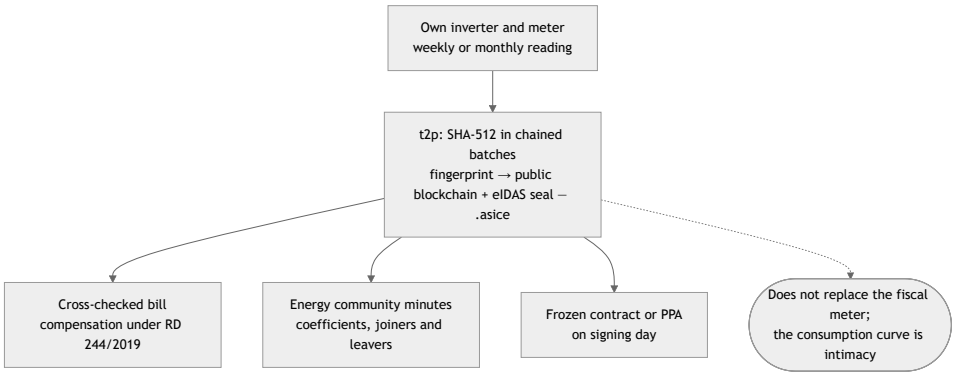


Figure 24.1 — Self-consumption with proof of existence (t2p: SHA-512 fingerprint on the device, RFC 3161/eIDAS, public blockchain, .asice): the inverter's and meter's own reading sealed in chained batches, the bill with its compensation cross-checked, the energy community's minutes (coefficients, joiners and leavers with a certain date) and the contract/PPA frozen on signing day. Context: 9.3 GW of self-consumption in Spain (UNEF) — fifth source by capacity; RD 244/2019 (simplified monthly compensation ≤100 kW; collective up to 5 km). Limits: one's own reading contrasts the fiscal meter, it does not replace it; calibration first; only fingerprints on the chain (the consumption curve is intimacy). Sources: UNEF; BOE (RD 244/2019).

---

[1] UNEF — *El autoconsumo fotovoltaico en España alcanza los 9,3 GW instalados* (report on 2025: +1,139 MW in the year, −3.7% vs 2024; 36,330 new households +229 MW; commercial +176 MW; industrial +679 MW; fifth generation source by installed capacity, ahead of nuclear; PNIEC target 19 GW); UNEF — *El autoconsumo sigue avanzando, superando ya el umbral de los 8 GW* (2024: +1,182 MW, cumulative 8,137 MW). https://www.unef.es/es/comunicacion/comunicacion-post/el-autoconsumo-fotovoltaico-en-espana-alcanza-los-9-3-gw-instalados · https://www.unef.es/es/comunicacion/comunicacion-post/el-autoconsumo-sigue-avanzando-superando-ya-el-umbral-de-los-8gw-instalados · https://www.eseficiencia.es/2026/02/02/autoconsumo-fotovoltaico-alcanza-9-3-gw-potencia-instalada-espana-segun-datos-unef ↩

[2] BOE — *Royal Decree 244/2019, of April 5* (administrative, technical and economic conditions of self-consumption; simplified surplus compensation ≤100 kW, voluntary, on the monthly bill — net billing; individual and collective self-consumption; administrative registry); MITECO — *Frequently asked questions on self-consumption*; UNEF — *Real Decreto 244/2019*; Puentes de Muras — *RD 244/2019: novedades* (extension of the collective self-consumption radius to 5 km, plants up to 5 MW). https://www.boe.es/buscar/doc.php?id=BOE-A-2019-5089 · https://www.miteco.gob.es/en/energia/energia-electrica/electricidad/autoconsumo-electrico/preguntas-frecuentes-autoconsumo.html · https://autoconsumo.unef.es/real-decreto-244-2019/ · https://puentesdemuras.com/es/rd-244-2019-normativa-fotovoltaica-espana/ ↩

[3] t2p — product manual (certify any file in batches — only the fingerprint travels; screen and web-page capture with full image + source code; document scanner up to 50 pages; photo with forensic GPS; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the CAE, in chapter 23; the community minutes, in chapter 37; the invoice, in chapter 34. ↩

Mobility I: Fleets and Companies, Solid Evidence

"It is a capital mistake to theorize before one has data."

— Arthur Conan Doyle, "A Scandal in Bohemia" (1891)

A second-hand van crosses a European border. In the country of origin, the odometer read 230,000 kilometers; by the time it reaches the destination dealership, it reads 150,000. Along the way, someone plugged a small device into the vehicle's diagnostic port and, in under a minute, erased **eighty thousand kilometers**. The car looks younger, is worth more, and the buyer will pay on average 26% extra for a vehicle that is in fact far more worn [3]. The most disturbing part is that this manipulation, according to a study by the **European Parliament's** research service written by Enrico Pastori and Raffaele Vergnani, *Odometer manipulation in motor vehicles in the EU* (2018), "leaves no trace": the electronics rewrite the number and the proof of the fraud evaporates [1].

It is no isolated case. That same study puts the cost of odometer fraud in the European Union at between **€5,600 and €9,600 million a year**, and estimates that between **30% and 50% of cars sold across borders** carry manipulated mileage [1]. The European Transport Safety Council confirms it: up to half of cross-border second-hand vehicles have a rigged odometer [2].

So far, an old scam of the car market. But let us pause a moment, because the story takes a turn that changes the whole picture. Those same sensors — the odometer, the GPS, the fuel consumption, the charging log — have ceased to be merely a car's technical file. Today they are the **raw material of new economic assets**: carbon credits, energy savings certificates (CAE) and corporate emissions reports. And if the data feeding them can be erased in a minute without a trace, all these instruments inherit the same fragility. This chapter is about how company fleets can generate climate evidence no one can falsify.

The fleet as a factory of savings — and of evidence

From a climate standpoint, a company fleet is a mine of savings. Electrifying it, training drivers in efficient driving, optimizing routes or switching to low-resistance tires generates real emissions reductions that can be turned into value. In Spain, the CAE system's catalog includes five transport-specific measure sheets: replacing tires with more efficient ones (TRA010), fleet **telemetry and geopositioning** systems (TRA020), long-distance shared vehicles (TRA030), urban collaborative mobility (TRA040) and replacing a combustion vehicle with a pure electric one (TRA050) [5]. Each measure, remember, is worth one kilowatt-hour saved and accredited.

And the potential is enormous, because the technology already works. The **International Council on Clean Transportation** (ICCT), an independent clean-transport research organization, measured ninety-one heavy electric trucks in real-world use in Europe in August 2025 and found they consume on average **65% less energy** than an equivalent diesel, with real range even better than the manufacturers advertised [6]. As the ICCT sums it up, the problem is no longer the technology, but **measuring and certifying the savings reliably**. On top of this pressure comes regulation: under the European sustainability reporting directive (CSRD), transport chain emissions — the so-called Scope 3 — must be disclosed when material, and are subject to external verification. And to count them there is even a common audited language, the Smart Freight Centre's **GLEC Framework**, aligned with the ISO 14083 standard [11].

The result is a paradox: we have the vehicles, we have the measure sheets, we have the accounting frameworks and the legal obligation to report. What is missing is the humblest link of all — the guarantee that the **source data** is true.

Sensors that lie

Because the attack surface of a connected fleet is enormous, every open door can inflate a saving or mint a false credit. The odometer can be rolled back — as we have seen — and the **GPS can be spoofed**: instead of jamming the signal, the attacker injects a false one, so a vehicle "appears" along efficient routes it never traveled, feeding carbon credits or transport CAEs. The **CAN bus**, the internal network connecting the vehicle's sensors, does not authenticate who sends each message; an injection through the diagnostic port

or the telematics can make the car say whatever suits. **Fuel cards** are a fraud classic: in 2022 an investigation coordinated by Eurojust dismantled a network across Romania, Belgium and Austria that had defrauded €420,000 with cloned cards and phantom refuelings traveling across half of Europe [8].

This is not theory. The cybersecurity firm Upstream counted **494 cyberattacks** on vehicles and fleets during 2025, 92% of them remote and two out of three aimed at telematics and the cloud [7]. Even electric charging has its flank: the OCPP protocol governing charge points authenticates each session, but weak versions allow the messages between vehicle and station to be manipulated, falsifying the kilowatt-hours billed or credited.

It is this book's recurring problem: *garbage in, garbage on-chain*. However sophisticated the registry where the credit is finally recorded, if the reading enters already falsified it remains immutably false. No blockchain saves a rigged odometer.

The lesson of Belgium: sealing works

Fortunately, we already know what happens when the data is sealed at the origin, because there is a full-scale experiment that proves it. It is called **Car-Pass** and it is Belgian. Before its introduction, Belgium suffered between **sixty thousand and one hundred thousand odometer manipulations a year**. The system is simple: every time a car changes hands, the mileage is officially recorded, and this history follows the vehicle forever. The result was devastatingly effective. In 2022, in all of Belgium, only **1,344 cases** of fraud were detected — under 0.2% of second-hand cars — and fraud in imports fell from 2.4% in 2020 to 0.87% in 2022 [4].

The lesson is luminous: when the truth of a data point is fixed at the moment it is generated, and that fingerprint accompanies the asset forever, fraud collapses. But Car-Pass has two limitations that make it insufficient for the world of fleets and carbon credits. It is **national** — a car entering from another country arrives with an empty or suspect history — and it is **centralized**, one database under a single authority. For a fleet operating across Europe, and for credits sold at global scale, we need the same Car-Pass principle but **decentralized, cross-border and verifiable by anyone**. And that is precisely what the proof of existence provides.

The proof of existence for fleets

Let us apply the mechanics of **t2p** to a fleet's data. The raw reading — the odometer's kilometers, the kilowatt-hours of a charging session, the GPS trace of a trip, the CAN bus packet — receives a **SHA-512 fingerprint** computed on the telematics device itself, as close to the source as possible, before any server can touch it. That fingerprint is accompanied by **forensic GPS** and a **liveness check**, which bind position and moment to the data and expose a phantom route or a spoofed GPS when the seal is crossed against real geometry. A qualified **RFC 3161 timestamp** is added: and here lies a powerful legal detail, because under Articles 41 and 42 of the European eIDAS regulation, this seal carries a **legal presumption of accuracy** of the date and time. In a dispute, it is not the party presenting the evidence who must prove the date is right; it is the party challenging it who must prove it is false [12]. Finally, the fingerprint is anchored to the **public blockchain** and the whole package is closed into an **.asice container verifiable offline**, which a CSRD auditor, a CAE verifier or a carbon-credit body can check without depending on the issuer.

It is, in essence, **Car-Pass raised to the power of the blockchain**: an incontrovertible history of mileage and consumption, but decentralized and cross-border. Sealing the reading at every refueling, every charging session, every technical inspection creates a record that makes odometer rollback, phantom trips to inflate savings, manipulation of charging sessions or the inflation of a CAE baseline impossible. And it dovetails naturally with two infrastructures already on their way. The first is the **Vehicle Identity** (VID) standard of **MOBI**, the *Mobility Open Blockchain Initiative*, a consortium founded in 2017 by manufacturers including BMW, Ford, General Motors, Honda and Renault, which in 2019 published the first digital "birth certificate" of a vehicle on blockchain and in 2024 demonstrated a global battery passport [9]. The second is the European **Battery Passport**, mandatory from February 18, 2027 for all electric vehicle batteries above two kilowatt-hours, which will record each battery's state of health, accessible through a QR code [10].

Honesty about the limit is due here too. The proof of existence guarantees that **the data has not changed since it was sealed**, not that the sensor told the truth at that instant. An odometer physically

rigged *before* the reading, or a sensor compromised in hardware, would still seal a false value. The mitigation is twofold: seal on the trusted device closest to the source and do it as often as possible, and cross-check redundant signals — GPS against consumption against kilometers — so that any incoherence leaps into view. The technology does not eradicate the lie; it makes it far harder, costlier and more traceable.

The promise and the peril

The promise, for a company with a fleet, is tangible. Evidence sealed at the origin turns carbon credits and transport CAEs into solid assets, defensible before an auditor and a buyer; it protects the company from unknowingly buying rigged vehicles; and it turns the Scope 3 reports the CSRD requires from an exercise of faith into checkable data. In a market where every tonne and every kilowatt-hour must be proven, reliable evidence is a competitive advantage.

The peril demands the usual nuance, the duality of "the promise and the peril." Telematics and the liveness check collect sensitive data — where each vehicle goes, who drives it — and demand a design scrupulous with workers' privacy. Concentrating verification in a few platforms would recreate the single point of trust we want to dissolve. And no seal exempts anyone from the physical work: measuring well, calibrating the sensors, auditing the baseline. The proof of existence is the foundation of trust, not a substitute for diligence.

It's time to act

For **fleet managers**: seal the key readings at the origin — mileage, consumption, charging sessions and routes — so that credits, CAEs and emissions reports are provable and unassailable. For **telematics providers**: build the cryptographic seal into the device, as close to the sensor as possible, not on a server where the data could already have been altered. For **CAE verifiers and carbon-credit registries**: demand anchored evidence as the basis of certification and use it to cut the cost of controls. For **sustainability auditors** (CSRD): require proofs of existence for Scope 3 data. For **regulators**: extend the Car-Pass logic — an unalterable mileage history — to a European, decentralized scale, and recognize the seal's evidentiary value, which your own eIDAS regulation already upholds.

A van can grow eighty thousand kilometers younger in one minute; but if the odometer reading is sealed at every change of hands, that minute ceases to exist. The difference between a fleet that claims to save and a fleet that **proves** it saves is, increasingly, the distinction between an asset and a suspicion. Generating that proof at the origin, tamper-proof, is time2act.

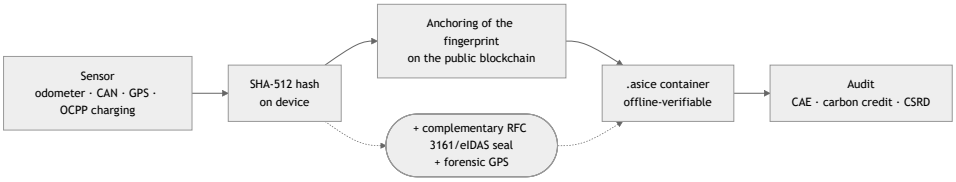


Figure 25.1 — A sensor reading (odometer, CAN bus, GPS, OCPP charging session) is digested into a SHA-512 fingerprint on the device, bound to forensic GPS and an RFC 3161/eIDAS timestamp with legal presumption, anchored to the public blockchain and packaged into an .asice container verifiable offline, ready to issue a CAE (sheets TRA010-TRA050), a carbon credit or a CSRD report. Comparison box: without a seal, up to 30-50% of cross-border odometers are tampered; with a sealed registry (the Car-Pass model), under 0.2%. Sources: European Parliament (EPRS); Car-Pass; MITECO.

[1] European Parliament (EPRS) — Enrico Pastori and Raffaele Vergnani, *Odometer manipulation in motor vehicles in the EU* (2018; damage of €5,600-9,600 million/year; 30-50% in cross-border sales). https://www.europarl.europa.eu/RegData/etudes/STUD/2018/615637/EPRS_STU(2018615637_EN.pdf ↩

[2] ETSC (European Transport Safety Council) — *EU should act against mileage tampering on second-hand cars* (up to 50% of cross-border cars). https://etsc.eu/eu-should-act-against-mileage-tampering-on-second-hand-cars/ ↩

[3] carVertical — *The cost of odometer fraud in Europe* (2024; ~€5.3 billion/year; average overprice of 26.3%). https://www.carvertical.com/en/blog/cost-of-odometer-fraud-in-europe ↩

[4] Car-Pass (Belgium) — *Annual report 2022* (1,344 cases in 2022, <0.2%; imports 2.4%→0.87%). https://www.car-pass.be/en/blog/car-pass-annual-report-2022 ↩

[5] MITECO — *CAE measure sheet catalog: transport sector* (sheets TRA010 to TRA050). https://www.miteco.gob.es/en/energia/eficiencia/cae/catalogo-de-fichas/catalogo-vigente-de-fichas/transporte.html ↩

[6] ICCT — *Real-world use cases for zero-emission trucks: heavy tractor-trailers for goods transport in the EU* (August 2025; 65% less energy; real range +11-19%). https://theicct.org/publication/rw-use-cases-for-zet-heavy-tractor-trailers-for-goods-transport-in-the-eu-aug25/ ↩

[7] Upstream Security — *2026 Global Automotive & Smart Mobility Cybersecurity Report* (494 cyberattacks in 2025; 92% remote; 67% on telematics/cloud). https://upstream.auto/press-releases/upstream-security-research-ransomware-attacks-on-automotive-and-smart-mobility-more-than-doubled-in-2025/ ↩

[8] Eurojust — *Suspects arrested in Romania, Belgium and Austria for defrauding EUR 420,000 using cloned fuel cards* (2022). https://www.eurojust.europa.eu/news/suspects-arrested-romania-belgium-and-austria-defrauding-eur-420000-using-cloned-fuel-cards ↩

[9] MOBI (Mobility Open Blockchain Initiative) — *Vehicle Identity (VID) standard* (2019) and *Global Battery Passport MVP* (2024). https://dlt.mobi/mobi-announces-the-first-vehicle-identity-vid-standard-on-blockchain-in-collaboration-with-groupe-renault-ford-and-bmw-among-others/ ↩

[10] Regulation (EU) 2023/1542 — *digital Battery Passport* (mandatory from 18-Feb-2027 for EV and industrial batteries above 2 kWh; QR access). https://www.circularise.com/blogs/eu-battery-passport-regulation-requirements ↩

[11] Smart Freight Centre — *GLEC Framework v3.0 / ISO 14083:2023* (freight transport emissions accounting). https://www.smartfreightcentre.org/en/about-sfc/news/a-solid-foundation-to-further-accelerate-freight-decarbonization-smart-freight-centre-releases-updated-glec-framework-version-30/ ↩

[12] Regulation (EU) 910/2014 (eIDAS), Articles 41-42 — legal presumption of accuracy of the qualified timestamp. https://eur-lex.europa.eu/eli/reg/2014/910/oj ↩

Mobility II: Shared Cars and Private Individuals, Distributing Trust

"Whoever you are... I have always depended on the kindness of strangers."

— Tennessee Williams, A Streetcar Named Desire (1947)

In Brooklyn, in August 2024, the federal prosecutor's office for the Eastern District of New York filed charges against two men over a scam of unsettling elegance. Eliahou Paldiel and Carlos Arturo Suárez Palacios had programmed a **GPS-spoofing** app and distributed it to more than eight hundred drivers on a ride-hailing platform. With that tool, a driver could "appear" in a high-fare zone — the so-called *surge* areas — without actually being there. Multiplied across hundreds of drivers and thousands of rides, the deception came to defraud some **forty million dollars** since November 2018 [5]. A second app, called "Screwber," fed them the destination and fare of each request in advance so they would accept only the most profitable trips.

The case is almost a parable. If the proof that a trip existed is a **blue dot on a map**, that dot can be forged. And the entire economy of shared mobility — from the car two neighbors share to a city's bike system — rests, ultimately, on this fragile proof. In the previous chapter we saw how a company with a fleet can seal the evidence of its vehicles. Now we descend to the hardest scale of all: that of the **private individual** who shares a car and who, with no central operator to vouch for them, wants to prove that a trip really happened in order to claim their fair share of a saving, a credit or a cost. Here, distributing trust is the central problem.

The quiet revolution of sharing

Sharing a vehicle is one of the most powerful and least glamorous climate levers there is. The reason is as simple as it is devastating: a private car sits parked about 95% of the time. This is stressed by **Susan Shaheen**, professor of engineering at the University of California, Berkeley and co-director of its Transportation Sustainability Research Center, probably the world's leading researcher on shared mobility. In a three-year study of the car2go service, Shaheen and Elliot Martin found that **each shared vehicle removed between seven and eleven private cars** from the streets of five North American cities, and reduced its users' miles driven and emissions [1].

That figure, however, demands an honest caution that runs through the entire literature: **the number of private cars one shared vehicle replaces is always a range, never a constant**. A European study of the SHARE NOW service, with more than 278,000 users across eleven cities, found an average of **3.3 cars replaced** — from 2.1 in Madrid and Lisbon to 5.3 in Brussels — while other cities reach far higher figures depending on the methodology [2]. Anyone selling you a single round number (the classic "one shared car replaces ten") is selling marketing, not science.

The BlaBlaCar platform avoided **2.5 million tonnes of CO_2** in 2024, filling more than seventy million empty seats and carrying ninety-two million passengers, as we have seen [3]. In North America, 2023 saw **157 million trips** on shared bikes and scooters, an all-time record [4]. It is a universe in full expansion, loaded with climate potential — so much so that peer-to-peer car-sharing platforms like Turo have gone as far as certifying themselves "carbon neutral" [11] — and it generates certifiable savings: carbon credits, collaborative mobility CAEs — recall sheet TRA040 from the previous chapter — or simple fair cost-splitting, provided one can prove the trip existed.

The problem of proof between equals

And here we hit a difficulty that corporate fleets do not have. When a company seals the data of its trucks, there is a responsible organization behind it. But when two strangers share a trip, **who attests that it happened?** Today, the answer is the app: BlaBlaCar, Turo or whoever acts as judge and party at once, and its "proof" is a record inside its own database, perfectly editable. To split a credit between two private individuals who do not know or trust each other, we depend blindly on a central intermediary — exactly the dependence this book has questioned since its first chapter.

This fragility opens the door to a specific, cheap kind of fraud. **Phantom trips** can be invented to "farm" credits that were never generated. Two accounts can **collude** to trade fake trips. And, as the Brooklyn case

illustrates, a **spoofed GPS** can fabricate entire routes that never existed. In every case, the attack is not against the savings formula but against the proof that the physical activity happened. As long as that proof is an editable blue dot, any incentive scheme for shared mobility is born with a hole in its heart.

A new currency: mobility credits

That hole is precisely what has held back one of the most fascinating ideas in transport economics: **tradable mobility credits**. The concept was formalized by the engineers **Hai Yang and Xiaolei Wang** in the article *Managing network mobility with tradable credits* (2011). The proposal is as elegant as it is equitable: instead of charging tolls, an authority would distribute to every citizen a quota of mobility credits, charge for use of the network in credits and let them trade freely. Whoever travels little — often, whoever has least — could **sell** their surplus to whoever travels a lot, achieving the same goals as a toll but with revenue neutrality and far more social justice [6].

The idea has spawned a thriving line of research. **Oded Cats**, professor of transport at Delft University of Technology, has studied the design and classification of these schemes and has even modeled their application to long-distance travel in Europe [7]. And yet there is a revealing fact: **as of 2025, no tradable mobility credit scheme has been implemented at real scale**. A recent academic review of the state of the art confirms it bluntly: only laboratory experiments, simulation games and modeling exist [8]. The incentives that do work today are more modest: in Belgium, the corporate "mobility budget" already reaches some 18,000 workers who trade the company car for sustainable alternatives [9]; at Delft University, citizens "spend" simulated credits in a game room; and in Brussels, the Region offers between €505 and €1,010 a year to those who deregister their car's license plate in exchange for public transport, bikes or shared cars [10] — but all of this is incentives, not yet credit markets.

Why does the idea never quite take off? Privacy, the legal framework and public acceptance all weigh. But there is a technical piece often forgotten, and it is exactly this book's: **without a reliable way for every citizen to prove their trips without handing their whole life to a central authority, the system is neither verifiable nor acceptable**. A mobility credit market needs, before anything else, a layer of trustworthy proof.

Distributing trust without a central operator

This is where the **proof of existence** of **t2p** supplies exactly the missing piece, and does so with a property that sets this chapter apart from the rest of the book. Let us apply it to a shared trip. On closing the trip, the individual's app generates **on their own device** a structured record — forensic GPS route, occupants, distance, mode — and computes a **SHA-512** fingerprint, so the raw data never leaves the phone: only its fingerprint does. That fingerprint receives an **RFC 3161 timestamp** conforming to eIDAS and is anchored to the **public blockchain**, so no one can later "create" a past trip to inflate credits.

And then comes the keystone, the piece the shared-mobility sector needs more than any other: the **proof of life**. Liveness biometrics tie the seal to the **real, present person** on that trip — not to an account or an emulator. This strikes at the root of the Brooklyn scam: an emulator can fake coordinates, but it cannot pass a liveness challenge anchored to the fingerprint's capture at that instant and that place. Finally, the whole bundle is closed into an **.asice container verifiable offline**, and here is the decisive turn: **two strangers can settle their shared trip without trusting any operator**, because both can verify the same proof independently. Trust is no longer lent by a central app; it is lent by mathematics.

This layer is, in reality, the infrastructure that was missing for the mobility credits of Yang and Cats to stop being a laboratory experiment. It allows credits to be distributed "without a trusted operator," prevents phantom trips and collusion between accounts, and makes any retroactive rewriting of the record detectable. It also fits open vehicle-identity standards like the VID of MOBI, the *Mobility Open Blockchain Initiative* [12].

Once again, honesty about the limit is due. The proof of existence certifies that *this record existed, intact, at that moment and tied to this person*; it does **not** guarantee that the underlying GPS reading is physically true if the phone's hardware or firmware was compromised before sealing. A GPS spoofed at the hardware level could still pass false coordinates. The real defense is the combination of layers — proof of life, coherence among altitude, speed and route, and the timestamp — not the blockchain alone. Nor does it prove **additionality**: that the shared trip really replaced a car that would genuinely have been used.

That remains a methodological assumption. The technology raises confidence in the evidence; it does not invent the truth of the physical world.

The promise and the peril

The promise is large and concrete for this sector: a peer-to-peer proof layer could finally unlock mobility credit markets, make cost-splitting between individuals fair and verifiable, and turn millions of small daily gestures — sharing a car, taking the bike — into accounted, rewarded climate savings without a big intermediary taking a cut.

But here the peril is especially sharp, and it has a name: **privacy**. What is being sealed are the movements of specific people. An infrastructure that proves where every citizen has gone can be, if badly designed, the most perfect surveillance tool ever built. That is why the property that the hash is computed on the device and the raw data never leaves it is not a technical detail: it is the line separating a tool of empowerment from a tool of control. The proof of life involves sensitive biometric data. And a badly conceived credit market could punish those with no alternative to the car — the resident of a village without a train — instead of helping them. The proof of existence must be, on this terrain, scrupulously privacy-respecting and equitably designed, or it will be nothing good.

It's time to act

For the **shared-mobility platforms** (BlaBlaCar, Turo, bike and scooter operators): offer users proofs of existence of their trips that are theirs and portable, not prisoners of the company database. For **private individuals**: demand the ability to prove your own trips without surrendering your entire location history to a third party. For the **designers of mobility credit schemes** (transport authorities, researchers like Cats): build the peer-to-peer proof layer in as a foundation, so your models can finally leave the laboratory. For **cities and regulators**: build these systems with privacy as the first requirement, not a final patch.

Two men in Brooklyn proved that a blue dot on a map is no proof at all. Rebuilding trust at mobility's most intimate scale — that of the citizen who shares a ride and wants it to count — demands a proof that neither an emulator nor an interested intermediary can fabricate. Distributing that trust, person to person, is the hardest and most human frontier in this whole book. And achieving it, too, is time2act.

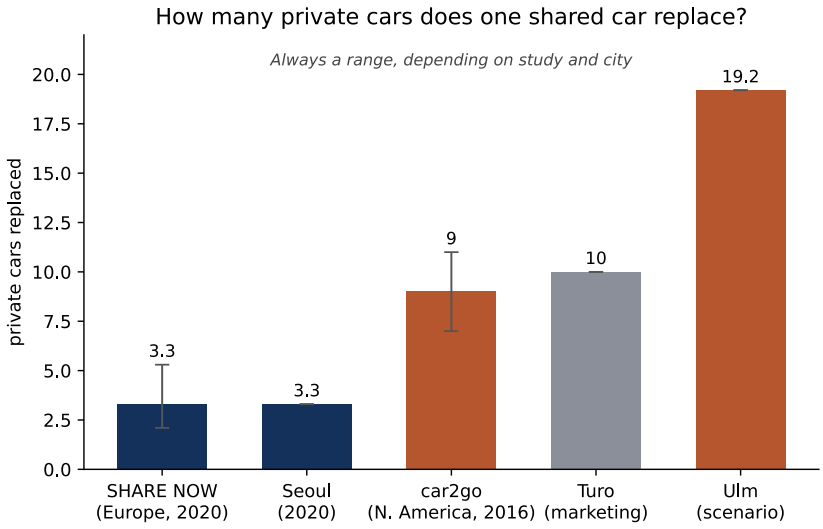


Figure 26.1 — The number of private cars a single shared vehicle removes from the street varies enormously by city and methodology: from an average of 3.3 in the European SHARE NOW study to the 7-11 of the car2go study in North America, with even higher scenarios in some cities. The publicity figure "one for ten" is commercial, not consensus. The lesson: never a single number; always a range with its source. Sources: UC Berkeley TSRC (Shaheen and Martin, 2016); Transportation Research Part A (2020).

[1] Susan Shaheen and Elliot Martin (UC Berkeley, Transportation Sustainability Research Center) — *Impacts of car2go: One-Way Carsharing in North America* (2016; each shared vehicle removes 7-11 private cars; VMT and GHG down). https://phys.org/news/2016-07-car-mobility-decreases-greenhouse-gas.html ↩

[2] *Does free-floating carsharing reduce private vehicle ownership? The case of SHARE NOW in European cities* (Transportation Research Part A, 2020; average of 3.3 cars replaced; range 2.1-5.3). https://pmc.ncbi.nlm.nih.gov/articles/PMC7544607/ ↩

[3] BlaBlaCar — *Sustainability Report 2024* (2.5 Mt CO_2 avoided in 2024; 92 million passengers). https://newsroom.blablacar.com/news/sustainability-report-2024 ↩

[4] NACTO — *A micromobility record: 157 million trips on bike and scooter share in 2023*. https://nacto.org/latest/a-micromobility-record-157-million-trips-on-bike-and-scooter-share-in-2023/ ↩

[5] Insurance Journal (EDNY prosecutors) — *Two charged in $40M ride-hail surge-pricing GPS spoofing scheme* (the "Fake GPS" case, August 2024). https://www.insurancejournal.com/news/east/2024/08/30/790683.htm ↩

[6] Hai Yang and Xiaolei Wang — *Managing network mobility with tradable credits* (Transportation Research Part B, 45(3), 2011, pp. 580-594). https://ideas.repec.org/a/eee/transb/v45y2011i3p580-594.html ↩

[7] Oded Cats et al. — *Design and classification of tradable mobility credit schemes* (Transport Policy, 136, 2023) and *Tradable mobility credits for long-distance travel in Europe* (Transportation Research Part A, 186, 2024). https://ideas.repec.org/a/eee/trapol/v136y2023icp59-69.html ↩

[8] Aljamal et al. — *Tradable credit scheme for mobility management: state of the art* (Transportation Research Interdisciplinary Perspectives, 2025; no real-scale implementation to date). https://www.sciencedirect.com/science/article/pii/S2590198225002908 ↩

[9] Next Mobility / ONSS — *Mobility budget in Belgium: ONSS figures 2024* (18,386 workers with a mobility budget; 17,030 without a company car). https://www.nextmobility.be/en/post/mobility-budget-belgium-onss-figures-2024/ ↩

[10] Bruxelles Environnement — *Prime Bruxell'Air* (aid of €505-1,010/year for deregistering a license plate, in force in 2024). https://environnement.brussels/citoyen/services-et-demandes/primes-et-aides-financieres/la-prime-mobilite-bruxellair-pour-se-deplacer-autrement-dans-la-capitale ↩

[11] SCS Global Services / Turo — *Turo becomes first global carbon-neutral certified car-sharing marketplace* (2023). https://www.prnewswire.com/news-releases/turo-becomes-first-global-carbon-neutral-certified-car-sharing-marketplace-301799877.html ↩

[12] MOBI (Mobility Open Blockchain Initiative) — *Vehicle Identity (VID) standard on blockchain*. https://www.traffictechnologytoday.com/news/smart-cities/mobi-creates-the-first-vehicle-identity-standard-on-blockchain.html ↩

Transport Delays: Proving the Arrival Time

"It is not that we have a short time to live, but that we waste a great deal of it."

— Seneca, On the Shortness of Life, I (c. 49 AD)

Any platform, any boarding gate. The display jumps from "+10" to "+45" with no explanation; the connection is lost; the mood tightens. Months later, the reply to the claim arrives in impeccable prose: "denied due to operational circumstances beyond the company's control." And the traveler discovers the heart of the matter: he has the memory of the wait; the company has the record — and the record is kept by the party that would have to pay.

The context is topical [1]: on July 1, 2024, Renfe cut back its voluntary punctuality commitment for AVE, Avlo, Alvia, Euromed and Intercity trains. Compensation for delays of 15 and 30 minutes disappeared, and the new regime pays 50% of the ticket from 60 minutes and 100% from 90 — while also incentivizing payment in points (200%) or vouchers (150%) instead of euros. The official justification: adapting to liberalization and European rules. The debate over whether compensation should start at 15 minutes remains open [1]. And the evidentiary lesson is immediate: when the thresholds rise, the frontier between 58 and 61 minutes is worth half a fare — and the official measurement of that frontier is made by one of the parties.

This chapter deals with the clearest and worst-evidenced consumer right there is. The claim to the company's customer service is in chapter 33; the cloud provider's clock — the same pattern in another industry — in chapter 59; fleets and shared mobility, in chapters 25 and 26. The key question: whose clock decides whether you get paid? The answer: the clock of whoever has to pay — unless the real arrival has a certain date on the traveler's device, with **t2p**: fingerprint, certain time and public anchor — the machinery of chapter 13 [3]. And with the usual symmetry: the carrier's anchored record protects it from the inflated claim.

The right exists; the proof, not so much

The framework for flights is among the most generous in European consumer law [2]: Regulation (EC) 261/2004 sets compensation of €250 for flights up to 1,500 kilometers, €400 for longer intra-European flights and those of 1,500 to 3,500, and €600 for those over 3,500; European case law — the Sturgeon doctrine — equated a delay of three hours or more at arrival with cancellation; and the Court of Justice has left a precious clarification for this book: the "arrival time" is the moment the doors open — an instant the passenger witnesses with a phone in their pocket. The exception that strains everything: the "extraordinary circumstances" that exempt the airline, and which the traveler can hardly cross-check.

The proof mechanism, by contrast, belongs to another age. The real arrival time is recorded by the carrier's systems; the display boards update and erase themselves; the reason for the delay is announced by the party that has to justify it; and the claim is filed weeks later, with a boarding pass and hazy memories. On the train, the same with the Spanish twist [1]: the commitment belongs to the house, and so does the measurement.

The conclusion asks no one to distrust anyone: it asks that minute 61 be a fact and not a version. The right already exists; what is missing is that its only variable — how late you arrived — can be attested from both sides.

The proof of existence applied

The arrival, turned into a fact [3].

1. The arrival with place and time

The decisive gesture takes ten seconds: as the doors open — the Court of Justice's definition [2] — a photo with t2p's forensic GPS: genuine fix, accuracy of about 50 meters, maximum age of about 30 seconds; a mocked location blocks certification [3]. Certain place and time of the real arrival, on the passenger's device. And if the flight was in airplane mode, no problem: background certification resumes on its own when coverage returns, and the seal attests the moment of capture [3]. Minute 61 stops being an argument.

2. The frozen display board

The board with the announced delay — or the alleged reason — the app screen with the flight or train status, the cancellation email: captured and certified with a date [3]. When the official version changes — the delay "for technical reasons" that later becomes an "extraordinary circumstance" — what was said at each moment exists, exact and dated, and no one can rewrite it after the fact.

3. The trip file

The ticket and the boarding pass, scanned; the expenses the company should have covered and did not — the meal, the hotel, the taxi — with receipts photographed with forensic GPS [3]; everything grouped into a project and exportable as an .asice package verifiable offline [3]. The claim to the company, to the state aviation safety agency or to the consumer court arrives with dated facts, not with a story.

4. The carrier's symmetry

The same mechanism protects the good operator: the record of departures and arrivals anchored in batches — the probe logic of chapter 59 — and the evidence of the genuinely extraordinary circumstance — the airspace closure, the storm, the third-party incident — sealed as it happens [3]. Against the inflated claim, or the passenger who never flew, the carrier with verifiable records answers with dates; and its punctuality commitment, anchored and public by fingerprint, becomes a commercial argument no competitor can imitate without delivering [1].

The limits, without hedging: the seal attests times, places and documents — it does not decide whether a circumstance was extraordinary: that is for the Regulation and the judge to say [2]. It does not replace the formal claim within the deadline or each authority's procedure. And proportionality rules: you document your own trip — never crews, employees or other passengers.

The promise and the peril

The promise: claims resolved on certain times [2][3]; passengers collecting what European law grants them without battles of versions; diligent carriers attesting their circumstances and settling legitimate claims in days; and a market — now liberalized [1] — where verifiable punctuality competes.

Now, the perils.

The first is **the industrialized claim**. Fraud has two faces: there are also passengers who claim what is not owed. Precisely for this reason the dated file benefits everyone: whoever flew and arrived late can prove it; whoever did not cannot fabricate a forensic GPS that refuses to certify mocked locations [3].

The second is **the war of clocks**. The passenger with their photo against the company with its systems. The way out is chapter 59's: official operating records anchored and public by fingerprint as the industry standard — when the official clock is verifiable, the private one is needed only by exception.

The third is **the trip as a profile**. Where you are and when you are there is intimacy pure and simple. The trip file stays on the device; on the chain, only fingerprints [3]; you disclose the leg being claimed, never your history of movements.

The fourth peril is not technical but one of perspective: mistaking the proof for the service. No seal makes the train arrive on time or prevents the storm. The proof makes the compensation enforceable and the punctuality statistics honest; the service is improved by investment and management.

It's time to act

Let the delay's clock have two owners.

For you, traveler: the photo at the opening of the doors [2][3] — ten seconds worth up to €600 [2]; and keep the display board and the expenses in the same project.

For the airlines and rail companies: anchor your operating records and publish punctuality by fingerprint [1][3] — the unfounded claim will die of dates, and your reliability will be verifiable, not promotional.

For the claims platforms: adopt the passenger's sealed file as the standard intake [3]; you will win cases with facts, not with volume.

For the authorities — aviation and rail: require operators to keep verifiable arrival records [1][2]; traveler protection cannot depend on the clock of the party that has to pay.

For legislators: the debate over thresholds — the 15, the 60, the 90 minutes [1] — is legitimate; what is not legitimate is that the measurement of the delay be unilateral. Let every new threshold be born with a verifiable arrival time.

The next time the board jumps from "+10" to "+45," let your clock count too — with a certain date. For every minute that is yours, **time2act**.

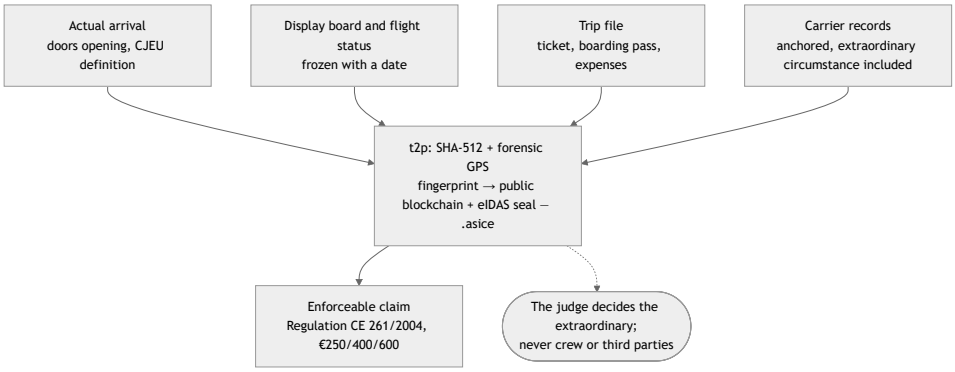


Figure 27.1 — The delay with proof of existence (sealed with t2p — chapter 13): the arrival with forensic GPS as the doors open (the CJEU's definition), the display board and flight status frozen, the trip file (ticket, boarding pass, care expenses) and the carrier's symmetry (anchored operating records; the extraordinary circumstance sealed as it happens). Context: Regulation (EC) 261/2004 (€250/400/600; Sturgeon ≥3 h); Renfe 1-7-2024: from 15/30 min to 60/90 (50%/100%). Limits: the judge decides what is extraordinary; document your own trip, never third parties. Sources: EUR-Lex; Renfe/Newtral.

---

[1] Newtral — *Cómo cambió Renfe las indemnizaciones por retraso del AVE* (1-7-2024: compensation at 15 and 30 minutes eliminated for AVE, Avlo, Alvia, Euromed and Intercity; new regime 50% from 60 min and 100% from 90; 200% refund in points or 150% in vouchers; justification: liberalization and legal alignment); Renfe — *Compromiso de puntualidad* and press note on the update; Xataka — *the debate over compensating from 15 minutes* (open regulatory situation). https://www.newtral.es/renfe-indemnizacion-retraso-ave-gobiernos-pp-psoe/20250801/ · https://www.renfe.com/es/es/ayuda/compromiso-puntualidad · https://www.renfe.com/es/es/grupo-renfe/comunicacion/renfe-al-dia/sala-de-prensa/renfe-actualiza-compromiso-voluntario-puntualidad-ave-avlo-alvia-euromed-intercity · https://www.xataka.com/movilidad/renfe-esta-obligada-a-indemnizar-retrasos-15-minutos-a-partir-1-enero-gobierno-quiere-impedirlo ↩

[2] EUR-Lex — *Regulation (EC) 261/2004* (compensation of €250/400/600 for cancellation and, under the CJEU's Sturgeon doctrine, long delay ≥3 hours at arrival; extraordinary circumstances exception; right to care; CJEU case law: the "arrival time" is the opening of the doors). https://eur-lex.europa.eu/legal-content/ES/TXT/?uri=CELEX:32004R0261 ↩

[3] t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; background certification that resumes when coverage returns — airplane mode; certified screen capture; document scanner; projects to group a trip's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the services claim, in chapter 33; the provider's clock, in chapter 59. ↩

Insurance: Proving the Claim

"Someone must have slandered Josef K., for one morning, without having done anything wrong, he was arrested."

— Franz Kafka, The Trial (1925), opening sentence

A British self-employed tradesman had, like everyone else, photos of his van on social media: the tool of his trade, parked in front of a building site, gleaming and without a scratch. In 2024, someone took one of those public images, digitally edited a crack onto the bumper and used it to file a fraudulent insurance claim, with a fake repair invoice for more than a thousand pounds, over an accident that had never happened. The fraud team at the insurer LV= uncovered it because the image was identical to the one on social media, but with damage added. Matt Crabtree, head of financial crime intelligence at LV=, warned that fraudsters use technology "for their illegal ends and to target innocent citizens for profit" [6].

The case is small in amount, but immense in what it announces. For decades, faking a claim required a complicit garage, an accommodating doctor or a genuinely dented car. Today an app on a phone is enough. And the proof on which the entire insurance system rests — the photograph of the damage — has suddenly become the easiest link to fabricate. This chapter is about how to make that photograph credible again: by sealing it at the moment and in the place it is taken, before the lie can get into it.

A theft we all pay for

Insurance fraud is not a victimless crime; we all pay for it through higher premiums. Its magnitude is dizzying. In the United States, the **Coalition Against Insurance Fraud** (CAIF), a coalition of insurers, regulators and consumers, published in 2022 — with Colorado State University Global — the first updated calculation in twenty-seven years: fraud costs the country **$308.6 billion a year**, the equivalent of some $933 for every American [1]. Health takes the biggest share (about $105 billion), followed by life and by property and liability.

In Europe, the insurers' federation **Insurance Europe** estimates that fraud — detected and undetected — amounts to around **10% of total claims spending**, a figure carried over from a 2013 report and still repeated because it is as robust as it is hard to improve [2]. The national records confirm it year after year. In the United Kingdom, insurers detected **£1.16 billion** in fraudulent claims in 2024 — 2% more than the year before — spread across 98,400 claims, of which more than half were motor [3]. In Spain, the sector detected **285,407 suspicious cases** in 2024 and prevented 63.7% of them; motor concentrated 61.2% of the fraud, with an average amount per prevented fraudulent claim of €2,924 and an investigation cost of barely €43 per case — a return of some €68 for every euro invested in investigating [4]. According to the barometer of the insurer Línea Directa, 8.8% of motor accident claims in Spain over the previous two years were attempted scams, almost five times more than in 2009 [10].

The new frontier: the claim that never happened

These figures were already high before artificial intelligence. Now, generative AI has put a factory of false evidence in everyone's pocket, and the sector is living it as an earthquake. The risk-data analyst **Verisk**, together with the National Insurance Crime Bureau, published in March 2026 a study — a survey of a thousand consumers and three hundred claims professionals — with blood-chilling data: **36% of American consumers** would consider digitally altering an image or a document in a claim even if it broke the insurer's rules, and the figure soars among the young — **55% of Generation Z**, 49% of millennials, against 12% of older people [5]. 13% even find it acceptable to "create photos of damage that never happened." And from the other side, **99% of insurers** have already encountered documentation manipulated or altered with AI, and 98% say these tools are increasing digital fraud [5].

This is not theory. The insurers Allianz and Zurich documented a **300% increase** in claims with manipulated photos comparing 2021-22 with 2022-23, and Allianz UK warned that the phenomenon had "all the hallmarks of becoming the sector's next big scam" [6]; in the first half of 2025 alone, Allianz uncovered in the UK more than **15,800 cases** of fraud worth **£92.6 million** [7]. Fraudsters buy photos of wrecked cars at online auctions, insert a real license plate with AI and fabricate collision damage. Zurich's team has caught many "by observing anomalies in the image metadata and with error level analysis," a forensic technique that betrays editing; in one case, the photo of the vehicle carried

timestamps predating the date of the supposed accident [8]. The reinsurer **Swiss Re**, in its *SONAR 2025* report, cites a **700% increase** in deepfake incidents in the fintech sector in 2023 and explicitly recommends combining biometrics, **liveness detection** and **unalterable marks anchored on blockchain** [9].

Detecting after the fact is not enough

The sector's immediate response is detection: analyzing every incoming image for traces of editing — pixel anomalies, incoherent metadata, error levels. It is valuable, necessary work, but it has a structural defect this book has already argued for deepfakes: it is an **arms race** in which the defender always runs behind. Every improvement in the detector pushes the fraudster toward better AI, and the day the generated image leaves no pixel trace, detection will go blind. Analyzing after the fact, moreover, means trusting metadata that the phone itself writes and anyone can rewrite.

The underlying problem is that the proof — the photo of the damage — is generated **outside** the insurer's control and **after** the event. The claimant sends a JPEG through an app or an email, and that JPEG can be a stock photo, an image from a scrapyard, pre-existing damage or an AI creation. As long as the proof is born like that, without anchoring, the entire architecture of the claim rests on sand. The lasting solution is not to analyze the incoming photo better, but to **guarantee the photo's origin at the moment it is taken**.

The proof of existence: sealing the claim at the scene

This is where the **proof of existence** of **t2p** attacks the problem at the root. Instead of examining the image once received, it seals it at the instant of capture, at the scene of the damage. The photo or video is **digested into a SHA-512 fingerprint on the device itself** at the moment it is taken; any later re-editing — adding a crack, altering a license plate — would change the fingerprint and break the chain. To the fingerprint are added an **RFC 3161 timestamp** conforming to eIDAS and an anchor on the **public blockchain** that fix *when* it was captured. And here is the key: this neutralizes exactly the trick that gave away Zurich's fraudsters — metadata with dates years earlier — because the seal comes from a trusted authority, not from the phone's manipulable EXIF.

Two more pieces are added that this sector needs especially. **Forensic GPS** attests *where* the photo was taken, which helps rule out auction images or damage recycled from elsewhere. And the **proof of life** ties the capture to *a real person present at the scene*, making it very hard for someone to upload images generated "from behind a computer," the pattern Allianz and LV= describe. All of it — the intact original image, fingerprint, seal, GPS, proof of life — is packaged into an **.asice container verifiable offline**, which the adjuster or the auditor can validate without depending on the app or a server, and which is admissible as evidence. It is worth underlining: this is precisely the recipe Swiss Re recommends — biometrics, liveness, unalterable anchoring — turned into a product.

As always, honesty about the limit is due. The proof of existence demonstrates that *this* image existed intact, in *this* place, at *this* moment, captured by *this* living person. **It does not demonstrate the cause or the legitimacy of the damage.** Someone can photograph with t2p after staging an accident themselves: the seal will be impeccable and the claim will still be fraud. That is why the proof of existence **complements, not replaces, the adjuster**: it removes the layer of documentary and image fraud — the fastest-growing one — so that human investigation can concentrate on causation and simulation. And it only adds value if the capture is made with the app at the moment, not for claims that already happened and were documented with old photos.

The promise and the peril

The promise is tangible and quantifiable. In a sector where fraud costs hundreds of billions, eliminating photo-evidence fraud at the root would free enormous resources, speed up payment of honest claims — the vast majority — and ease the pressure on everyone's premiums. A photo sealed at the scene turns the most vulnerable part of the process into the most solid.

The peril calls for the usual nuance, the duality of "the promise and the peril." Collecting geolocated photos and liveness biometrics touches sensitive client data at a moment of stress; the design must be scrupulous with privacy and must not turn every claim into an excessive surrender of data. No seal

exempts the adjuster from judging the cause, nor the insurer from treating the client with trust by default. And concentrating verification in a few platforms would recreate the single point of trust we want to dissolve. The proof of existence is a necessary layer within a broader architecture of trust.

It's time to act

For insurers: offer origin-sealed capture as the preferred channel for the claim report and reward it with faster processing — whoever has nothing to hide comes out ahead. For honest policyholders (the vast majority): demand the ability to prove your claim with evidence no algorithm can cast doubt on. For adjusters and investigators: reserve human effort for causation and hard fraud, letting the technology close the door on image fraud. For regulators: recognize the seal's evidentiary value and set standards that protect the integrity of the evidence and the client's privacy at the same time.

A tradesman watched the photo of his van become the weapon of a scam. The difference between an image anyone can fabricate and a proof sealed at the scene is, increasingly, the key to a system that defends itself. Making the claim photo credible again — proving the damage existed, there and then — is, here too, time2act.

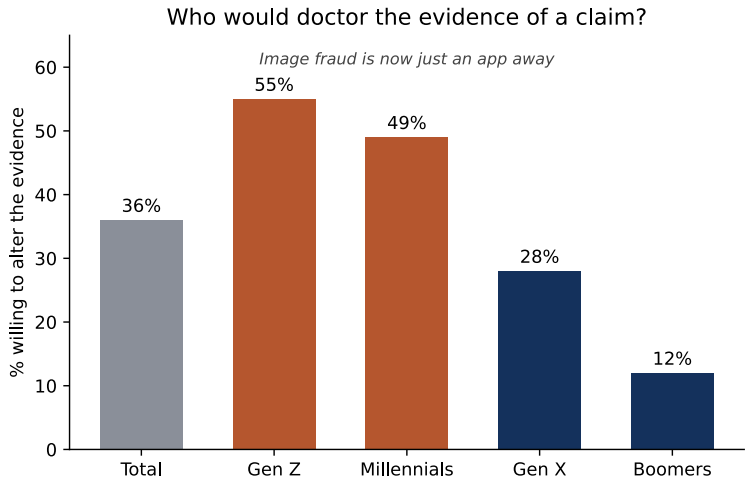


Figure 28.1 — Percentage of consumers willing to digitally alter the image or document of a claim, even against the insurer's rules: 36% on average, but up to 55% among Generation Z. The willingness falls with age. This is the face of the AI image fraud that 99% of insurers have already seen. Source: Verisk with the National Insurance Crime Bureau (2026).

---

[1] Coalition Against Insurance Fraud — *The Impact of Insurance Fraud on the U.S. Economy* (2022; $308.6 billion/year, ~$933 per person). https://insurancefraud.org/wp-content/uploads/The-Impact-of-Insurance-Fraud-on-the-U.S.-Economy-Report-2022-8.26.2022.pdf ↩

[2] Insurance Europe — *Annual Report 2023-2024* (Fraud chapter) and *The impact of insurance fraud* (~10% of claims spending; 2013 estimate, ~€13 billion referring to 2017). https://www.insuranceeurope.eu/downloads/ar-2024-fraud/Fraud.pdf ↩

[3] Association of British Insurers (ABI) — *Fraudulent insurance claims continue to top £1 billion* (2024: £1.16 billion; 98,400 claims). https://www.abi.org.uk/news/news-articles/2025/11/fraudulent-insurance-claims-continue-to-top-1-billion/ ↩

[4] ICEA — *El sector detectó 285.407 casos sospechosos de fraude, de los que evitó un 63,7 %* (Spain, 2024). https://segurosnews.com/news/el-sector-detecto-285-407-casos-sospechosos-de-fraude-de-los-que-evito-un-637 ↩

[5] Verisk (with the National Insurance Crime Bureau) — *AI editing tools are fueling a new era of insurance fraud* (March 2026; 36% of consumers, Gen Z 55%; 99% of insurers). https://www.verisk.com/company/newsroom/ai-editing-tools-are-fueling-a-new-era-of-insurance-fraud-according-to-new-research-from-verisk/ ↩

[6] The Guardian (via PetaPixel) — *Fraudsters are editing photos of damaged vehicles to claim insurance* (LV= case; +300% manipulated photos, Allianz/Zurich). https://petapixel.com/2024/05/07/fraudsters-are-editing-photos-of-damaged-vehicles-to-claim-insurance/ ↩

[7] Allianz UK — *Allianz UK detects frauds totalling £93m in first half of 2025* (15,800 cases). https://www.allianz.co.uk/news-and-insight/news/allianz-uk-detects-frauds-totalling-93m-in-first-half-of-2025.html ↩

[8] Zurich Insurance — *AI poses challenges, offers tantalizing solutions to insurers fighting fraud* (metadata + error level analysis). https://www.zurich.com/insights/business/ai-poses-challenges-offers-tantalizing-solutions-to-insurers-fighting-fraud

↩

[9] Swiss Re Institute — *SONAR 2025: How deepfakes, disinformation and AI amplify insurance fraud* (+700%; recommends liveness and blockchain anchoring). https://www.swissre.com/institute/research/sonar/sonar2025/how-deepfakes-disinformation-ai-amplify-insurance-fraud.html ↩

[10] Línea Directa — *VII Barómetro del Fraude al Seguro* (8.8% of motor claims in Spain). https://www.lineadirectaaseguradora.com/-/vii-barometro-de-fraude-al-seguro ↩

Supply Chain: Verifiable Origin Against Counterfeiting

"It is not from the benevolence of the butcher, the brewer, or the baker, that we expect our dinner, but from their regard to their own interest."

— Adam Smith, The Wealth of Nations, I, 2 (1776)

In the autumn of 2022, in the hospitals of Banjul, the capital of The Gambia, doctors began to see babies arriving with an acute kidney failure they could not explain. The trail led to four pediatric cough-and-cold syrups, manufactured by an Indian company, Maiden Pharmaceuticals, and contaminated with diethylene glycol and ethylene glycol — two toxic industrial solvents. Close to **seventy children** died. On October 5, 2022, the World Health Organization issued Medical Product Alert No. 6/2022 [3]. Three years later, in October 2025, the story repeated itself in India: the "Coldrif" syrup, with nearly 48% diethylene glycol, killed at least twenty-four more children, and the WHO had to issue a new alert [4].

The detail that changes everything is this: in those cases, **the label was authentic and the documentary chain was "legitimate."** The papers were in order. What failed was not the seal or the barcode, but the physical truth of the contents. It is the hardest, most honest lesson of this chapter, and it deserves to come first: no proof of origin, however perfect, analyzes the molecule. But, as we shall see, that does not make the proof of origin useless — it makes it indispensable for everything else in the problem.

The fraud that enters legitimate chains

Product counterfeiting is no longer confined to brand-name bags sold on a street blanket; it is a river infiltrating legitimate supply chains and threatening health. The **OECD** and the **European Union Intellectual Property Office** (EUIPO), in their report *Mapping Global Trade in Fakes 2025*, put world trade in counterfeit and pirated products at **$467 billion** in 2021 — **2.3% of all the world's imports** [1]. In the European Union the share rises to **4.7% of imports**, some €117 billion [1]; as we have seen, what alarms most is the shift toward **dangerous counterfeiting**: fake car parts, cosmetics, toys, food and medicines entering the market as if they were genuine.

Medicines are the most lethal case. The **World Health Organization** estimates that **at least one in ten medical products** in low- and middle-income countries is substandard or falsified, based on more than a hundred studies and 48,000 samples from 88 countries [2]. The economic waste hovers around **$30.5 billion** a year, but the real cost is counted in lives like The Gambia's. The WHO distinguishes the **substandard** product — one that fails quality standards through poor manufacturing — from the **falsified** one — one that deliberately misrepresents its identity, composition or origin. In both cases, the fundamental problem is that **there is no easy way to prove where what we hold in our hands really comes from**.

The regulator demands coordinates

Faced with this, Europe has taken a historic turn: it now demands that origin be *verifiable*, not merely declared. The most ambitious piece is the **Deforestation Regulation**, the EUDR (Regulation EU 2023/1115). For the first time, the law requires the **geolocation coordinates of the exact plot** from which each consignment of seven commodities comes — cattle, cocoa, coffee, palm oil, rubber, soy and wood — to guarantee they do not come from land deforested after December 31, 2020 [5]. The mechanics are revealingly concrete: for plots over four hectares, a **polygon** of coordinates drawing its perimeter is required; for smaller ones, a single point with six decimal places suffices; everything is submitted in GeoJSON format to the Union's information system. The application dates are set at **December 30, 2026** for large and medium operators and **June 30, 2027** for the smallest [6]. It is, in one sentence, the shift of proof "from paper to land."

It is not the only example. The **Falsified Medicines Directive**, in force since February 2019, requires most prescription medicines to carry a **unique identifier** — a 2D code with product, serial number, batch and expiry — and an anti-tampering device, verified against a central European repository before dispensing [7]. And in the world of luxury, the **Aura Blockchain Consortium**, as we have seen, has already registered more than **fifty million products** on a blockchain that gives the buyer direct access to

the history and proof of authenticity of their piece [8][9]. Three different responses, one and the same intuition: trust must be provable.

Why the seal alone does not save (and why it is needed anyway)

At this point, one must be implacably honest, because this chapter touches lives. The serialization of a medicine proves that **the box** is authentic, not that **the molecule** inside is. In The Gambia, a sealed photo of the factory would not have detected the toxic solvent. It is the principle that runs through the whole book, here in its most literal form: "garbage in, garbage on-chain." A proof of existence of the **wrong** plot, or of a correct batch with adulterated contents, is still a well-sealed lie.

Why, then, is it worth it? Because the vast majority of supply chain fraud is **not** deliberate poisoning, but the **backdating** of certificates, the **substitution of documents**, the falsification of origin and the reuse of photos from another plot or another year. Against all of that, sealing the act of origin is devastating: it raises the cost of lying, fixes responsibilities with a name and a date, and leaves a trail an auditor can follow to the source. The proof of existence does not replace physical quality control — the kind that would have saved the children of The Gambia; it **complements** it, closing the door on documentary fraud so that inspection can concentrate on what truly matters: the contents.

The proof of existence at the origin

Let us apply the mechanics of **t2p**. At the point of origin — the plot, the harvest, the factory, the batch — a georeferenced photograph is taken and sealed on the spot: a **SHA-512 fingerprint** computed on the device, **forensic GPS** fixing the coordinates, an **RFC 3161 timestamp** conforming to eIDAS and an anchor on the **public blockchain**. For the EUDR, this means photographing the plot with its polygon or its six-decimal point and turning the capture into **unalterable evidence** of the exact moment and place. The **proof of life** can tie the evidence to the person conducting the inspection or the harvest — the inspector, the farmer — reducing fraud because "no one can seal this without being there." And the whole bundle is packaged into an **.asice container verifiable offline**, which a customs officer, an auditor or a buyer can check without a connection and without trusting the supplier's word.

Most importantly, the proof of existence **does not compete** with serialization or digital passports; it complements them. The medicine's 2D code and Aura's passport say *what* product it is; t2p's origin seal adds the missing layer: *this batch was photographed here, on this date, with this person present*. It is the difference between having a trust label and having the **proof of the act** that the label summarizes. For the luxury consumer who scans their bag and sees its history, and for the customs officer verifying a container of coffee, the promise is the same: being able to check the origin without having to believe anyone.

The promise and the peril

The promise is enormous and transversal. A supply chain with its origin sealed would be far more resistant to documentary forgery, would meet the EUDR with real evidence instead of declarations, would protect honest brands and, above all, would give citizens the right to check what they buy. In a trade where counterfeits move close to half a trillion dollars, recovering trust in origin is a competitive advantage and a public health matter.

The peril calls for the usual nuance, in the duality of "the promise and the peril." The first, already explained: the seal certifies the act, not the physical quality; presenting it as a guarantee of safety would be dangerous and false. The second, the burden that can fall on small producers of the Global South, for whom geolocating every plot can be a barrier; the design must be inclusive, not a new wall. And the third, privacy and concentration: whoever controls the verification layer must not become a new single intermediary. The proof of existence is a necessary tool within a broader architecture that always includes physical control.

It's time to act

For **producers and farmers**: seal the act of origin — plot, harvest, batch — so that traceability is demonstrable and defensible before the EUDR and the buyers. For **importers and brands**: demand

anchored evidence as the basis of due diligence; do not accept backdated certificates. For **regulators and customs**: recognize the seal's evidentiary value and integrate it into the information systems, never forgetting the physical control of the product. For **consumers**: be able to scan and check the origin of what you eat, take or wear. And for everyone, remember the lesson of The Gambia: verifiable origin is necessary, but product safety also demands analyzing its contents.

Seventy children died with an authentic label in hand. No technology should claim the power to avert every tragedy, but we can make lying about origin much harder and more traceable. Proving where something really comes from — and who takes responsibility — is, here too, time2act.

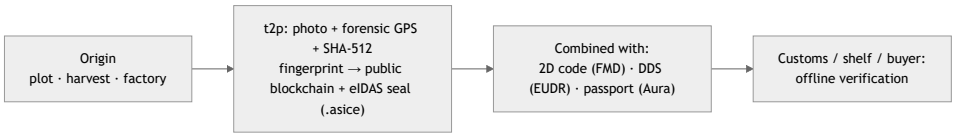


Figure 29.1 — At the point of origin (plot, harvest, factory) a georeferenced photo is captured and sealed with forensic GPS, an eIDAS timestamp and a SHA-512 fingerprint, and packaged into an .asice container verifiable offline. This proof-of-the-act-of-origin layer combines with serialization (the FMD's 2D code), the EUDR due diligence declaration or Aura's digital passport, through to final verification at customs or on the shelf. Sources: EUDR (EU 2023/1115); EMA; Aura Blockchain Consortium.

---

[1] OECD / EUIPO — *Mapping Global Trade in Fakes 2025* (counterfeits: $467 billion in 2021, 2.3% of world imports; EU 4.7%). https://www.oecd.org/en/publications/mapping-global-trade-in-fakes-2025_94d3b29f-en.html ↩

[2] World Health Organization — *Substandard and falsified medical products* (≥1 in 10 in low/middle-income countries; ~$30.5 billion/year). https://www.who.int/news-room/fact-sheets/detail/substandard-and-falsified-medical-products ↩

[3] WHO — *Medical Product Alert No. 6/2022* (The Gambia; Maiden Pharmaceuticals syrups contaminated with diethylene glycol). https://www.who.int/news/item/05-10-2022-medical-product-alert-n-6-2022-substandard-(contaminated-paediatric-medicines ↩

[4] WHO — *Medical Product Alert No. 5/2025* (India; "Coldrif" syrup by Sresan Pharmaceuticals). https://www.who.int/news/item/13-10-2025-medical-product-alert-n-5-2025--substandard-(contaminated-oral-liquid-medicines ↩

[5] EUR-Lex — *Regulation (EU) 2023/1115* (EUDR; plot geolocation coordinates; 7 commodities; cutoff date 31-Dec-2020). https://eur-lex.europa.eu/eli/reg/2023/1115/oj/eng ↩

[6] European Commission (Access2Markets) — *Delay until December 2026 — implementation of the EUDR* (second postponement, Reg. EU 2025/2650; application 30-Dec-2026 / 30-Jun-2027). https://trade.ec.europa.eu/access-to-markets/en/news/delay-until-december-2026-and-other-developments-implementation-eudr-regulation ↩

[7] EMA — *Falsified medicines: overview* (Directive 2011/62/EU; mandatory 2D unique identifier since 9-Feb-2019). https://www.ema.europa.eu/en/human-regulatory-overview/public-health-threats/falsified-medicines-overview ↩

[8] Aura Blockchain Consortium — *About* (founded in 2021 by LVMH, Prada and Cartier; >50 million products registered). https://auraconsortium.com/about ↩

[9] WWD — *Aura Blockchain crosses 50 million products* (LVMH, Prada, OTB; September 2024). https://wwd.com/business-news/technology/aura-blockchain-crosses-50-million-products-lvmh-prada-otb-1236606994/ ↩

Consumer Purchases: Fake Reviews, Inflated Prices and Warranties on the Clock

"Publicity is about social relations, not objects. Its promise is not of pleasure, but of happiness: happiness as judged from the outside by others."

— John Berger, Ways of Seeing (1972)

Before buying anything, you do what everyone does: you read the reviews. Five stars, "it changed my life," thousands of ratings that sound like the collective voice of a satisfied crowd. What you may not know is that behind that unanimity there can be an industry. In July 2022, Amazon sued the administrators of more than 10,000 Facebook groups dedicated to orchestrating fake reviews in exchange for free products or money: a single group, "Amazon Product Review," had more than 43,000 members; in another, $10 was paid per review; the scheme operated on the stores of the United States, the United Kingdom, Germany, France, Italy, Spain and Japan, and dodged Meta's detection by swapping letters in sensitive words [2]. Since 2020, Amazon had reported more than 10,000 groups to Meta; only half had been removed [2].

This chapter deals with the layer of information on which the consumer decides. Where the product physically comes from is explained in chapter 29; who pays for the ads nobody sees, in chapter 56; what is really inside the food, in chapter 101. Here the questions are the shopping cart's: can I believe these stars? Is this discount real? And if it breaks, will I be able to prove when and how? The book's answer: give the certain date back to the three signals — the review with an attested purchase, the sealed price series and the warranty chronology — with **t2p**: the SHA-512 fingerprint on the device, the anchoring of the fingerprint to the public blockchain and the complementary RFC 3161/eIDAS timestamp, inside an .asice container [5]. Neither the market nor the consumer should ever again argue over memories.

The fabricated opinion

Amazon's lawsuit put numbers to an open secret: opinions are mass-produced and sold like any other product. Organized groups with tens of thousands of members, per-review rates, evasion techniques to fool the platforms' filters [2] — the complete professionalization of a fraud that decides purchases every second.

The regulatory response has arrived. In the United States, the FTC's final rule — announced on August 14, 2024, in force since October 21, 2024 — bans buying or selling fake reviews, "insider" reviews without disclosure of the relationship, "independent" review sites controlled by the company itself, the selective suppression of negative reviews and the sale of fake social-influence indicators, with penalties of up to $51,744 per violation [1].

But the law, on its own, does not solve the epistemological problem: the reader still cannot tell, star by star, who has touched the product and who is paid to opine. And with generative AI, the cost of fabricating a thousand plausible opinions — with human touches, invented usage details and the tone of a sincere neighbor — tends to zero.

The way out is not to believe less: it is to be able to verify more. If every review can be tied to a sealed proof of purchase — the receipt fingerprinted on the buyer's device, without revealing their identity [5] — the platform can distinguish, and the reader can filter, the opinion of someone who actually bought. You can buy a thousand opinions; you cannot fabricate a thousand sealed purchases before writing them. The star stops being an act of faith.

The discount that wasn't

The second fabricated signal is the "before" price. The OCU has measured it in its Black Friday analyses, and the picture is bleak: only one in three products actually gets cheaper; in categories like smartwatches, cameras and headphones, prices rise on average 13% in the weeks before the campaign; 43% of products spike in price during the campaign itself; the average advertised saving is 23%, but real prices come out on average 3% more expensive [3]. And the regulatory datum: in 99% of cases, the indicated "previous price" is not the lowest of the last 30 days, as the sales rules require — many retailers use as reference a "manufacturer's recommended price" that has often never been the real selling price [3].

The OCU can detect this because it does what no consumer can do alone: monitor prices from early September [3]. And this is exactly where the proof of existence changes the game: the rule exists, but its verification is inaccessible to the citizen — you would have had to watch the price every day.

The anchored price series does it for everyone: if prices are recorded with sealed captures — every price, every day, with a certain date [5] — by the platform itself, by consumer observatories or by organized citizens, the "previous price" becomes a fact verifiable in one click, and the fake discount a provable infringement. The "−70%" sign stops being a declaration of faith: either it withstands the series, or it exposes whoever hangs it. And for the honest retailer, the same seal is marketing — the verifiable discount is worth more than the inflated one next door.

The warranty on the clock

The third signal is the promise that, if it fails, they will answer for it. The Spanish framework in force is among the most protective in Europe: since January 1, 2022 — transposing the European directives — the minimum legal warranty is three years for goods (two for digital content and services); during the first two years the defect is presumed to be of origin (one year for second-hand products); manufacturers must guarantee spare parts for ten years after the product leaves production; and the deadline to claim is five years [4].

In practice, however, the warranty is a race of proofs against the clock. When did you buy? — the receipt has vanished among the invoices. How did it arrive? — the scratch "was already there" or "you did it yourself." When did you report the defect? — inside the two-year presumption, or outside it. How many times has it been repaired? Every unproven question is a claim that dies before the after-sales desk.

The application of the proof of existence is direct. The receipt and the invoice, sealed on the day of purchase, fix with a certain date the start of the three years [4][5]: the lost paper stops being fatal. The product's condition at delivery — four photos or thirty seconds of video fingerprinted in the moment — is the consumer's equivalent of the rental inventory: neither the courier nor the buyer can rewrite how the box arrived. And every notice to the seller, anchored, fixes whether the defect appeared within the presumption [4]. The warranty stops being a battle of memory: no more fighting the clock — now the clock works for you.

The proof of existence applied

An ordinary purchase, three ten-second seals — the receipt, the opened box, the first notice — and no possible argument months later.

1. The review with an attested purchase

The real buyer seals the proof of purchase on their device and publishes the review accompanied by the proof: the fingerprint attesting the purchase exists, without revealing the buyer's name or details [5]. The platform marks these opinions with a verifiable badge; the reader filters by "attested purchase"; and the black market in reviews loses its raw material [2]. The FTC rule supplies the fines [1]; the proof of existence supplies the mechanism that makes them almost unnecessary.

2. The price series that does not lie

Consumer observatories, honest platforms or organized citizens record prices with sealed captures — every price, every day, with a certain date [5]. When the discount season arrives, the "previous price" is checked against the anchored series: if 99% of the signs fail the 30-day rule [3], let the series prove it, not the buyer's patience. Price transparency stops depending on the good faith of the sign.

3. The opened box with a certain date

The day the package arrives, thirty seconds of video or four photos sealed on the spot: the product's condition, the serial number, the contents of the box [5]. If the product arrived damaged, the proof exists from minute one and the refund brooks no delay; if the buyer claims damage they caused themselves, the absence of contemporaneous proof exposes them. E-commerce's daily war — transport damage, returns fraud — is settled on images with attested integrity, in both directions: the honest consumer gets paid sooner, and the cheat runs into cryptography.

4. The warranty clock

The sealed purchase starts the three years with a certain date [4][5]; the first defect notice, anchored, fixes whether the problem appeared within the two-year presumption [4]; and every successive repair is dated, so the product that strings together breakdowns under warranty — the candidate for replacement or refund — carries its own chronology. For the diligent seller, the same record is the defense against the out-of-time claim. Europe's longest warranty [4] is worth what its proof is worth — and now the proof is automatic.

The limits, stated clearly: the seal attests that the purchase, the photo or the notice existed on a date — it does not attest that the review is sincere (a real buyer can also be paid to opine), nor that the product is good, nor does it settle who is right about a technical defect: that is for experts, platforms and consumer courts. And the attested-purchase review must always preserve the buyer's anonymity: attest the purchase, never identify the person [5].

The promise and the peril

The promise: a market where the stars have a purchase behind them [1][2], the discounts withstand the 30-day rule [3] and the three-year warranty is exercised without hoarding receipts in a shoebox [4]. Consumer trust on a verifiable basis: cheaper for the honest, costlier for the cheat.

The perils, with the same clarity.

The first is **the buyer's anonymity**. Attesting the purchase cannot mean identifying the person. If the "attested purchase" badge demanded a name and a profile, we would have traded fake reviews for consumer surveillance. The right design: the cryptographic proof of the purchase travels; the identity, never [5].

The second is **censorship by verification**. If platforms published only reviews with attested purchases, they would silence legitimate voices — the gift received, the purchase made elsewhere, the service experience. The badge must filter and weight, not censor: free opinion for everyone, verifiable weight for whoever attests it.

The third is **the bought seal**. Fraud adapts: sellers who "sell" products at a token price to harvest attested purchases, price series sealed by ghost shops. The seal attests facts, not good faith — pattern surveillance (bulk purchases, impossible prices) remains the job of platforms and regulators [1][3].

And the fourth is **consumer techno-solutionism**. No fingerprint will make products better or prices fairer; that is done by competition, consumer inspection and the rules that already exist [1][3][4]. The seal only guarantees that the rules of the game can be verified. A market can be perfectly transparent and perfectly expensive.

It's time to act

The consumer decides on signals; let the signals be trustworthy again.

For you, the consumer: seal the receipt on the day of purchase, the box on the day of delivery and the first defect notice [4][5]; before a spectacular discount, remember the datum — only one in three is real [3]; and before a thousand unanimous stars, look for the attested purchase [2].

For commerce and review platforms: implement the anonymous attested-purchase review [1][2][5] and publish anchored price series — credibility is your product.

For honest retailers: the discount verifiable against the sealed series [3][5] is your best argument against the inflated sign next door.

For consumer authorities: enforce the 30-day rule against anchored price series [3] and accept the sealed chronology — purchase, unboxing, notices — as preferred evidence in warranty claims [4].

For regulators: the FTC rule marks the path [1]; complete it with the mechanism — a presumption of validity for the anonymous attested purchase and penalties for fabricated social indicators.

Ten dollars per review, 43,000 members in a single group [2]: consumer trust had become a commodity. Let it go back to what it was — an honest signal between strangers. For the everyday shopping cart,

time2act.

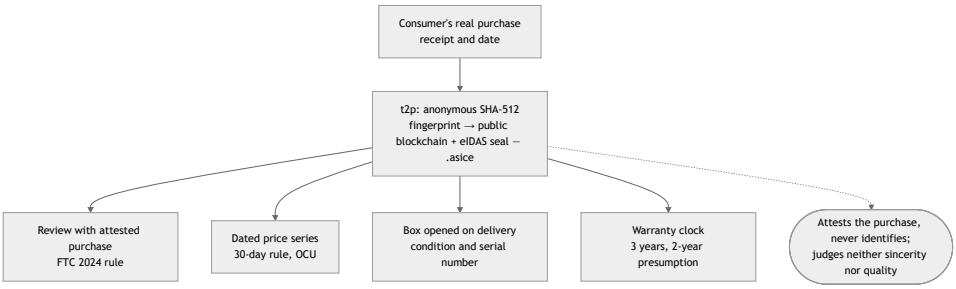


Figure 30.1 — The consumer's information layer with proof of existence: the review with an anonymous attested purchase (the receipt sealed on the device; the review black market loses its raw material), the price series with a certain date (the lowest-price-in-30-days rule becomes verifiable) and the warranty clock (purchase, unboxing and notices sealed activate and attest the 3 legal years). Limits: the seal judges neither sincerity nor quality; purchase attestation must never identify the buyer. Sources: FTC (2024 rule); Amazon vs. Facebook groups; OCU (Black Friday); CEC/Consumo (3-year warranty).

---

[1] Federal Trade Commission — *FTC Announces Final Rule Banning Fake Reviews and Testimonials* (August 14, 2024; in force October 21, 2024; up to $51,744 per violation; bans the sale of reviews, insider reviews without disclosure, controlled review sites, selective suppression and fake social-influence indicators); Federal Register — *Trade Regulation Rule on the Use of Consumer Reviews and Testimonials*. https://www.ftc.gov/news-events/news/press-releases/2024/08/federal-trade-commission-announces-final-rule-banning-fake-reviews-testimonials · https://www.federalregister.gov/documents/2024/08/22/2024-18519/trade-regulation-rule-on-the-use-of-consumer-reviews-and-testimonials ↩

[2] Infobae — *Amazon demanda a miles de grupos de Facebook por publicar reseñas falsas* (July 19, 2022; more than 10,000 groups; US, UK, Germany, France, Italy, Spain and Japan); Marketing4eCommerce — *Amazon denuncia a los administradores de más de 10.000 grupos de Facebook* ("Amazon Product Review," more than 43,000 members; $10 per review; more than 10,000 groups reported to Meta since 2020, half removed); La Nación — *Amazon demanda a los administradores de 10.000 grupos de Facebook*. https://www.infobae.com/america/agencias/2022/07/19/amazon-demanda-a-miles-de-grupos-de-facebook-por-publicar-resenas-falsas/ · https://marketing4ecommerce.net/amazon-reviews-falsas/ · https://www.lanacion.com.ar/tecnologia/amazon-demanda-a-los-administradores-de-10000-grupos-de-facebook-por-promover-resenas-falsas-nid19072022/ ↩

[3] OCU — *Black Friday: descuentos que no son reales* (only 1 in 3 products gets cheaper; average rises of 13% in the preceding weeks in categories like smartwatches, cameras and headphones; 43% of products spike in price; advertised saving of 23% vs real prices 3% dearer; in 99% of cases the "previous price" is not the lowest of the last 30 days; monitoring from September); Público — *La trampa del Black Friday: precios inflados, rebajas ficticias*. https://www.ocu.org/consumo-familia/compras-online/noticias/bf-precios-rebajados · https://www.publico.es/sociedad/trampa-black-friday-precios-inflados-rebajas-ficticias-fiebre-compras-detiene.html ↩

[4] European Consumer Centre Spain (Ministry of Consumer Affairs) — *Spain extends the minimum legal warranty from two to three years* (since January 1, 2022; transposition of Directives (EU) 2019/770 and 2019/771; presumption of origin of the defect extended from 6 months to 2 years — 1 year for second-hand; spare parts for 10 years; limitation period from 3 to 5 years); OCU — *Garantía de los productos: se amplía el plazo*. https://portal-cec.consumo.gob.es/en/comunicacion/noticias/2022/espana-amplia-la-garantia-legal-minima-de-dos-tres-anos · https://www.ocu.org/consumo-familia/derechos-consumidor/noticias/novedades-garantia ↩

[5] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; attestation without identification — the proof of purchase travels, the identity never; content on the device and voluntary disclosure). See chapters 13 and 15; the product's origin, in chapter 29. ↩

Second-Hand Between Individuals: The Item's Condition, With a Certain Date

"The maxim is: doveryai, no proveryai — trust, but verify."

— Ronald Reagan (with Gorbachev), signing of the INF Treaty (1987)

The buyer opens the box that has just arrived. Inside, where a games console should be, there is a brick wrapped in newspaper. At the other end of the country, a seller receives a return "for defects" and discovers that the phone sent back is not the one he shipped: a different casing, a screen with years of scratches. Two victims of the same void: between strangers, no one can prove the condition of an object at a given moment.

And the void is country-sized [1]: Wallapop alone has more than 19 million active users in Spain, and the second-hand market between individuals moves more than €13 billion a year. The dark side: scams on these platforms are the bulk of the more than 412,000 computer frauds recorded in 2024, complaints grew 60% in 2025, and the Guardia Civil issues periodic alerts about the campaigns [1].

This chapter deals with the economy's most informal contract: the sale between strangers. Fake reviews and the retail warranty are in chapter 30; the customer facing the big company, in chapter 33; the fake rental listing, in chapter 36; and the payment tricked outside the platform — recovering a transfer — is dealt with in chapter 69. The key question: how do you prove what condition an object was in when it changed hands? The answer: documenting it with a certain date at both ends of the journey — the video before shipping and the video on opening — with **t2p**: the triad of chapter 13 — fingerprint on the device, certain time and anchor on the chain [3]. The symmetry is total: the same gesture disarms the package that is not what was promised and the switcheroo of the return.

The symmetrical scams

The repertoire is documented [1], and it is worth facing it squarely.

The **off-platform payment**: the listing with impeccable photos and a price too good to be true, and the insistence on being paid by Bizum or bank transfer "to save on fees." Once the money is paid, nothing arrives — or anything at all does — and the buyer has lost the platform's protection along the way: the payment already lives outside its framework of guarantees. (The banking component — recovering the tricked transfer — is the subject of chapter 69.)

The **triangulation fraud and the "switcheroo"**: the scammer buys with shipping, receives the item in good condition and immediately files a complaint alleging defects — exploiting the protection that holds the payment until the buyer validates. He requests the return, keeps the good item and sends back a battered one. The honest seller is left without the money and without the product — and before the platform, he has only his word.

The **"confirm the payment" phishing**: the link that mimics the platform's interface and harvests credentials; and the **fraudulent shipping labels** and impersonations that dress the scam in legitimacy [1].

The evidentiary diagnosis: a sale between individuals has no invoice and no commercial warranty — the civil hidden-defect claim is the only substantive recourse — and all the evidence lives in listing photos without a reliable date, chats and ambiguous tracking. When the platform arbitrates, it decides with this: words and images anyone can fabricate or reuse. The honest stranger has no way to distinguish himself from the fraudulent stranger — and this is exactly the class of problem the certain date solves.

The proof of existence applied

Ten seconds at each end of the journey [3].

1. The item documented before shipping

The seller films the item with t2p before packing: its general condition, the serial number, checking that it works if applicable — with the date and time overlaid on the video and the forensic GPS of the place [3].

Thirty seconds that fix a fact: this object, in this condition, existed like this at this moment. The listing photo can be old or someone else's; the certified video cannot.

2. The sealing in a single take

The item going into the box, the closing, the shipping label being applied — one single shot, no cuts, certified [3]. It is the gesture that disarms the "there was nothing inside": the contents and the package are bound together in a single dated video. For items of value, it is the difference between a lost dispute and one closed in five minutes.

3. The symmetrical opening

The buyer performs the same gesture at the other end: the closed package with the label visible, the opening, the contents — in a certified single take [3]. If the brick is inside, the video attests it with date and place; if the item is as promised, likewise. And against the "switcheroo" [1], the pair of videos — the seller's sealing and the buyer's opening, and those of the return in the opposite direction — turns arbitration into a comparison: what left, what arrived, what was returned.

4. The dispute file

The listing captured with a date — the photos and description that were promised — the chat with the agreement, the payment receipt, the shipment tracking: grouped into a project [3] and exportable as an .asice package verifiable offline. For the platform's arbitration, for the complaint to the Guardia Civil [1], for the civil claim: dated facts instead of debatable screenshots.

And a first cousin of this economy between individuals deserves its own paragraph: the loan between family and friends. The "I lent it to you" that years later becomes "it was a gift" has broken more friendships than any scam. The brief agreement sealed at the moment the money changes hands — the amount, the repayment terms, two lines with a certain date [3] — protects the lender and, above all, protects the relationship: when the terms are a fact, memory does not have to play referee.

The limits, before anyone gets carried away: the seal attests conditions and moments — it does not decide whether the price was fair or whether the defect was a hidden vice: that is for the platform, the expert or the judge to assess. Proportionality rules: the €10 sale needs no file; the €500 one perhaps does — let each person calibrate. And you never film the courier or any third party: you document the object, not the people.

The promise and the peril

The promise: disputes resolved by comparing dated videos rather than by a moderator's intuition [1][3]; the "switcheroo" that stops being profitable because the returned item is checked against the sealing video; honest sellers and buyers who distinguish themselves from scammers with a ten-second gesture; and a €13 billion market [1] gaining the trust it lacked to grow.

The perils, without hedging.

The first is **evidentiary theater**. The scammer can film too: a video with the good item and the shipment of another. That is why method matters: the uncut single take through to the label, the forensic GPS that refuses mocked locations [3], and the symmetrical pair of videos — a performance has to fool both cameras and both dates at once, and that is no longer a cheap con.

The second is **friction that kills the sale**. If selling a second-hand book demands a film shoot, no one will do it. The answer is proportionality: the short gesture for ordinary sales, the full file for valuable ones — and platforms integrating it so it costs one click.

The third is **the platform as an opaque judge**. Internal arbitrations decide by their own rules with no real appeal. The verifiable trail makes it easy for them to decide well; the deeper step is for them to value anchored evidence above screenshots — and to say so in their rules.

And one last trap: trusting the tool to keep watch for you. No seal replaces elementary prudence: pay inside the platform, distrust Bizum transfers to strangers and "confirmation" links [1]. The certain date protects the legitimate transaction; the off-platform payment scam is fought by not taking the bait.

It's time to act

Let the honest stranger always win.

For sellers and buyers: the ten-second gesture at each end — the item, the sealing, the opening [3]; and for the sales that matter, the whole file.

For the platforms: integrate evidence with a certain date into your arbitrations [1][3] — you will resolve faster, expel the professionals of the "switcheroo," and your trust badge will be more than marketing.

For the carriers: the package's weight and condition at intake, anchored — your link in the chain counts too.

For the Guardia Civil and consumer authorities: complaints grow 60% a year [1]; the citizen's .asice file is the difference between shelving and investigating.

For legislators: buying and selling between individuals is the real economy of millions of people [1]; let digital consumer protection recognize anchored evidence — from both sides — as qualified proof.

The next package you send or receive, let it carry its moment of truth on record. For a neighborhood market where the honest stranger always wins, **time2act**.

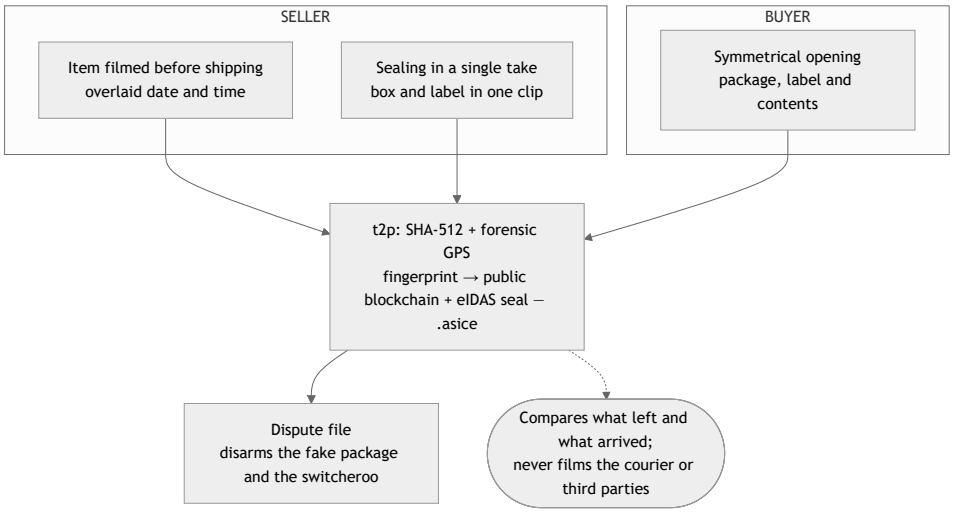


Figure 31.1 — Second-hand between individuals with proof of existence (sealed with t2p — chapter 13): the item filmed before shipping (overlaid date/time + forensic GPS), the sealing in a single take, the buyer's symmetrical opening and the dispute file (captured listing, chat, payment, tracking). Total symmetry: it disarms the fake package AND the "switcheroo." Context: >19 million Wallapop users; >€13 billion/year; second-hand scams as the bulk of the >412,000 computer frauds of 2024 (+60% complaints in 2025). Limits: proportionality by value; the hidden defect is for the expert; never film third parties. Sources: Guardia Civil/specialist press.

---

[1] Digital Perito — *Estafas en Wallapop y Vinted: cómo denunciar con pruebas* (>19 million active Wallapop users in Spain; second-hand market >€13 billion a year; second-hand platform scams as the bulk of the >412,000 computer frauds of 2024; +60% complaints in 2025); Público — *Así operan los estafadores en plataformas de segunda mano* (typologies: off-platform payment, triangulation/"switcheroo" exploiting buyer protection, phishing, fraudulent labels, impersonation); OKDiario / Genbeta — Guardia Civil alerts and detection tips. https://digitalperito.es/blog/estafas-wallapop-vinted-como-denunciar-pruebas-2026/ · https://www.publico.es/sociedad/asi-operan-estafadores-plataformas-segunda-mano-vinted-wallapop.html · https://okdiario.com/curiosidades/fraude-que-arrasa-plataformas-segunda-mano-guardia-civil-lanza-advertencia-sobre-vinted-wallapop-15912396 · https://www.genbeta.com/seguridad/estas-estafas-wallapop-comunes-varios-trucos-para-pillarlas-rapidamente ↩

[2] Civil framework of sales between individuals (hidden defects; absence of consumer commercial warranty). See chapter 30 for the B2C warranty and chapter 36 for the fake listing scam.

[3] t2p — product manual (video with overlaid date/time; photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; certified screen capture; certify any file; projects to group a sale's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the tricked payment, in chapter 69. ↩

Renting Things: The Condition at Pickup and at Return

"Neither a borrower nor a lender be, / For loan oft loses both itself and friend."

— William Shakespeare, Hamlet, I, 3 (c. 1600)

The airport, the rental-counter queue, the quick walk-around with the damage sheet in hand — "is this scratch already here?" — and the hurried return: the key drop-box, the flight. Three weeks later, the email: "damage charge: €460," with two photos taken who knows when.

It is not bad luck: it is a documented pattern [1]. The official report on irregular practices of Spanish rental companies describes disproportionate indemnities — €300 to €1,000 for small scratches or blemishes — charged directly to the deposit card, with the nature and the amount fixed unilaterally by the company; and the European Consumer Centre documents the later charges for cleaning or minor damage decided in inspections carried out without the consumer present, who is forced to bear them without express knowledge [1]. There have been convictions: one European leader of the sector was ordered to refund damages improperly charged [1].

And the irony that gives this chapter its cue: the ECC itself recommends returning the car while the office is open, checking it with the staff and "taking photographs of the interior and the exterior… they can serve as evidence" [1]. The regulator has already understood that this is about proof; what is missing is for the consumer's photos to be worth as much as the company's inspection.

This chapter deals with the contract over the condition of a thing. The bilateral inventory of the rented flat is in chapter 36; the condition of the item sold between individuals, in chapter 31; the peer-to-peer shared car, in chapter 26; the recurring charge and the claim to the company, in chapter 33; the card chargeback, in chapter 69. The key question: how do you argue about the condition of a thing at two moments if neither of the two moments has an opposable record? The answer: with the condition video at both ends of the rental — at pickup and at return — sealed with **t2p** — the usual machinery, that of chapter 13 [2]. And with the usual symmetry: the same trail protects the honest company from the customer who denies the dent.

The asymmetry of the counter

The problem, piece by piece. The **initial damage sheet**: a diagram with crosses that the customer signs in a hurry, often without light or time to review it. The **excess and the deposit**: the retained card that allows charging first and arguing later. The **return inspection**: when the office is closed or the customer is already flying, the condition is determined by the company alone, with no counterweight [1]. The **price of the blemish**: internal scales, not estimates — the scratch valued at hundreds of euros according to a table the customer has never seen [1]. And the **distance**: the foreign tourist who will not come back to claim.

The B2B extension reproduces the same scheme with bigger sums: the rental of construction machinery and professional equipment — the initial condition of the backhoe, the hours of use, the damage "discovered" at return — often between companies that will need each other next season.

The evidentiary diagnosis: a rental is a contract over the condition of a thing at two moments — and neither of the two moments has, today, a record that counts for both parties. The customer's photos have no opposable date; the company's inspection has no customer in front of it. When the charge arrives, the argument is the usual one: your word against their card.

The proof of existence applied

Two minutes at each end of the rental [2].

1. The pickup walk-around

Before leaving the parking lot, the complete walk-around filmed with t2p: the entire bodywork, the wheels, the interior, the dashboard with the kilometers and the fuel level — with the date and time overlaid and the forensic GPS of the parking lot (genuine fix, accuracy of about 50 meters; a mocked

location blocks certification) [2]. Every pre-existing scratch, documented; the sheet of crosses, scanned and sealed. Two minutes worth exactly whatever the company decides to claim later.

2. The return walk-around

The same gesture at return — and this is where everything is won or lost: in the after-hours return, with the key drop-box and no one from the company present [1], the certified video of the final condition is the only inspection with a certain date that will ever exist. The charge three weeks later [1] runs into a fact: this is how the car was on the day I returned it, at this hour, in this parking lot.

3. The contract and the charge

The contract, the condition inventory and the excess terms, scanned and sealed at signing [2]; and when the late charge arrives, the dated claim — to the company, to the consumer mediator, and if necessary the chargeback to the bank, the path chapter 69 explains. The .asice file — walk-arounds, contract, charge, claim — verifiable offline by anyone [2].

4. The honest lessor's symmetry

The company — or the individual renting out their machine — that documents with the customer present has the same protection: the handover inspection and the return inspection, filmed and sealed, with the customer present or notified [2]. Against the customer who denies the new dent, the pair of videos answers; and the damage price scale, published and anchored in advance, turns the price of the scratch into a known condition, not a surprise. In machinery B2B, this bilateral file is quite simply the difference between keeping a client and losing them in a lawsuit.

The limits, in black and white: the seal attests conditions and moments — it does not judge whether the blemish is normal wear or damage, or whether the price is proportionate: that is for the contract, the mediator and the judge [1]. The filmed walk-around must be complete and in good light — a dark video seals its own uselessness. And the usual proportionality: the ten-euro scooter needs no file; the car with a €1,200 excess does.

The promise and the peril

The promise: consumers returning the car after hours with the peace of mind of a dated fact [1][2]; late charges that die on sight of the file; honest companies settling real damage with bilateral evidence and setting themselves apart from the bad practices that soil the sector [1]; and the rental of things — from the holiday car to the backhoe — running on verifiable conditions.

And the perils, one by one.

The first is **the bad-faith video**: filming while dodging the dent you have already made. The method fights it: the complete walk-around in a single take, the overlaid date, the parking lot's GPS [2] — and the pair of videos: whoever hides a dent in their own walk-around reveals it in the other end's.

The second is **counter friction**. If picking up a car demands a ritual, the customer will not do it. Two minutes, a guided gesture — and the companies that build it into their own process (the joint filmed walk-around) will turn it into a selling point.

The third is **the unilateral price scale that survives**. Proof of condition does not fix the abusive price of the scratch [1]: that is the job of the consumer regulator and the courts. The certain date makes the dispute clean; the fairness of the scale is another battle.

There remains the fourth: the temptation to delegate everything to the seal. No seal replaces returning the car while the office is open when you can [1], or reading the excess before signing. The proof is for when prudence is not enough.

It's time to act

Let the condition of the thing have two witnesses: one at each end.

For you, the renter: the two-minute walk-around at both ends [2] — and if you return after hours, never skip it: it is the only inspection you will have.

For the rental companies: bilateral filmed inspections and price scales published by fingerprint [1][2] — the documented charge gets paid without argument, and the sector cleans itself of the reputation that follows it.

For professional and machinery rental: the handover-and-return file as a standard annex to the contract [2]; condition disputes are the most avoidable lawsuits in your business.

For the consumer authorities and the ECC: you already recommend the photos [1]; the next step is recognizing the consumer's evidence with a certain date as a formal counterweight to the unilateral inspection.

For legislators: the inspection without the customer present and the unilateral card charge [1] are the loophole; let every damage charge require verifiable condition evidence from both moments.

May the next airport key drop-box not be an act of faith. Two minutes of video, and the returned car stays returned — for both sides. For rentals without surprises, **time2act**.

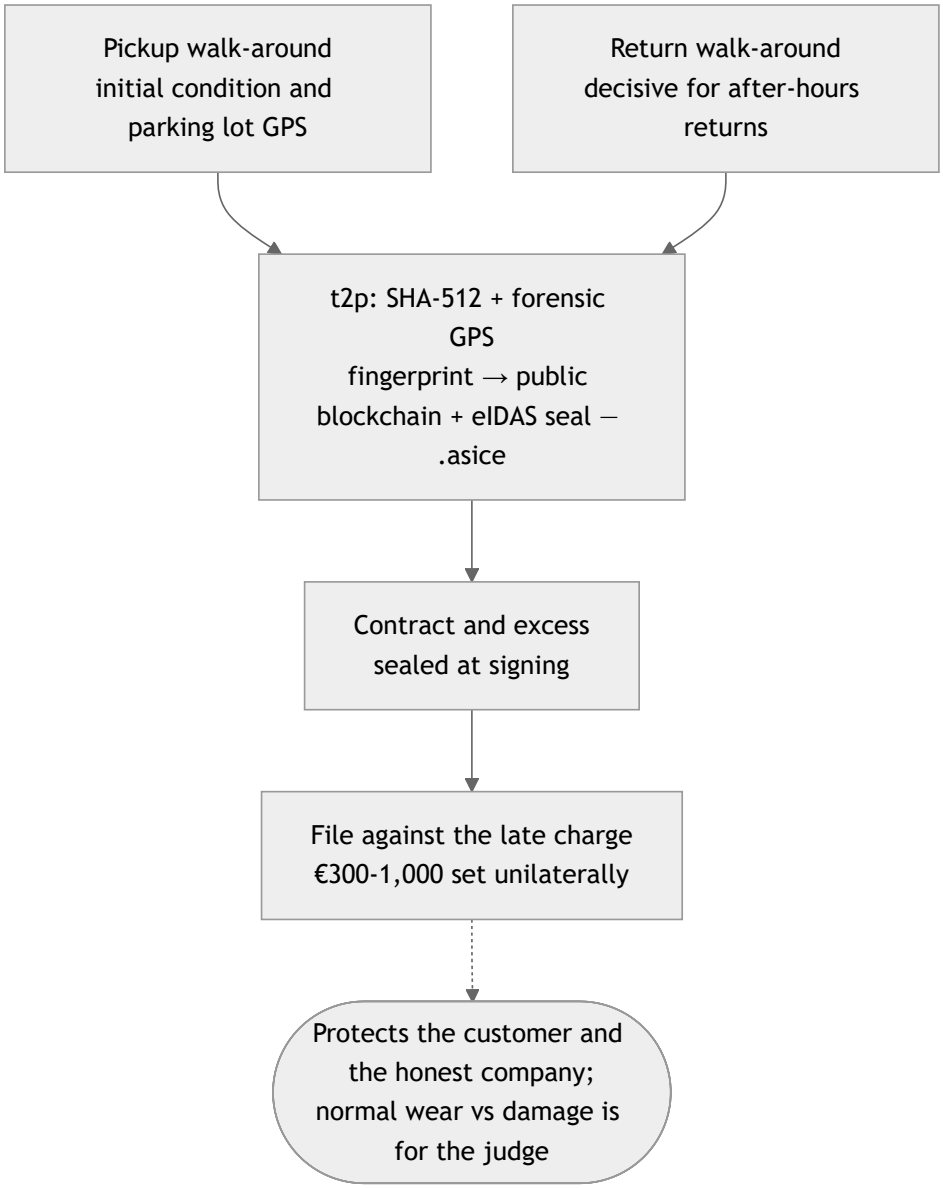


Figure 32.1 — Renting things with proof of existence (sealed with t2p — chapter 13): the pickup walk-around filmed (overlaid date/time + forensic GPS of the parking lot), the return walk-around — decisive in after-hours returns — the contract and the excess sealed, and the file against the late charge. Symmetry: the bilateral inspection protects the honest company from the customer who denies the dent; extendable to B2B machinery rental. Context: documented irregular practices (charges of €300-1,000 set unilaterally; inspections without the customer); the ECC already recommends photos — what was missing was the certain date. Limits: normal wear vs damage and prices are for the contract and the judge. Sources: Consumo/DSCA; CEC-España.

[1] Directorate-General for Consumer Affairs (DSCA) — *Prácticas irregulares de compañías españolas de alquiler de coches* (disproportionate indemnities of €300 to €1,000 for small scratches or blemishes, charged to the deposit card with the amount set unilaterally); CEC-España (European Consumer Centre) — *Alquiler de vehículos: derechos y recomendaciones* (later charges for cleaning or minor damage; inspections without the consumer present; recommendations: return with the office open, check with staff, request a receipt and "take photographs of the interior and the exterior as evidence"); AEA — *Líder europeo del rent-a-car condenado a devolver el importe de unos daños cobrados indebidamente*. https://www.dsca.gob.es/sites/default/files/consumo_masinfo/PRACTICAS_IRREGULARES_ALQUILER_COCHES.pdf · https://portal-cec.consumo.gob.es/sites/default/files/documentos/Folleto_Alquiler_Coches.pdf · https://portal-cec.consumo.gob.es/en/comunicacion/noticias/2021/cec-espana-recuerda-los-derechos-de-las-personas-consumidoras-cuando · https://aeaclub.org/rent-a-car-reclamacion-coche-alquiler-cobro-indebido/ ↩

[2] t2p — product manual (video with overlaid date/time; photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; document scanner; projects to group a rental's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the flat's inventory, in chapter 36; the sold item's condition, in chapter 31; the card charge, in chapter 69. ↩

The Customer Against the Company: Cancellation, Lock-In and the Call Center

"If the main pillar of the system is living a lie, then it is not surprising that the fundamental threat to it is living the truth."

— Václav Havel, The Power of the Powerless (1978)

The hold music, a long quarter of an hour, and at last a voice: "call back tomorrow, the system isn't working for me right now." The next day, another voice: "I'll transfer you to the retentions department." An offer, a promise — "we'll keep your price and the cancellation is on file" — and the feeling of being done. Three months later, the charge keeps arriving, and customer service replies with all the calm in the world: "no cancellation request is on record." Nobody remembers exactly what was said; all that remains is a verbal agreement that never materialized.

It is the country's most everyday consumer complaint, and the imbalance feeding it is structural: the company records every call "to guarantee service quality" — but the recording belongs to the company; the complaint code is generated by its system; the cancellation record lives on its servers. When the dispute comes, the customer brings memory; the company brings its database. Law 10/2025 has just given consumers new rights [1] — but a right without proof of one's own is a right that depends on the other side's good faith.

This chapter deals with the customer's evidentiary relationship with the big services company: the cancellation, the lock-in, the recurring charge, the call center promise. Reviews, prices and the product warranty are in chapter 30; the relationship with public administration, in chapter 44; the lawfulness and technique of recording one's own conversations — Constitutional Court ruling STC 114/1984 — in chapter 75; the cloned voice that calls pretending to be you, in chapter 68. The key question: how do you exercise a right when all the evidence lives in the adversary's house? The answer: documenting every transaction in the moment with **t2p** — the evidentiary package chapter 13 unpacks piece by piece [3] — so that the certain date finally sits on the customer's side too. And with the usual neutrality: the same trail vouches for the diligent company that serves well and processes on time.

New rights, old proof: Law 10/2025

The legislative novelty is first-rate [1]: Law 10/2025, of December 26, on customer service — published in the BOE on December 27, 2025, with twelve months for companies to adapt — imposes concrete, measurable obligations. Answer 95% of calls in under three minutes on average. Assign every inquiry, complaint, claim or incident an identifying code that allows its status to be tracked. Answer claims within fifteen days at most — and unauthorized charges within just five. Give fifteen days' notice of every automatic renewal, making cancellation easy. Make cancellation simple, without labyrinths: the ID number and the contract number suffice. And prohibit exclusively automated service: the customer has the right to a human operator.

On paper, it is a before-and-after. But look at it with this book's eyes and you will see the hole: the law grants rights; it does not grant proof. Who attests that the call lasted forty-seven minutes? That the cancellation was requested on the 3rd and not the 28th? That the charge was unauthorized and reported within the deadline? That the renewal arrived without notice? The tracking code is generated and held by the company; the entire file lives in its databases. A right without proof of one's own is exercised at the adversary's courtesy — and courtesy, in a call center with quarterly targets, is no procedural guarantee.

When the charge comes back: the cases

The sanctions exist, and they map the problem [2].

Vodafone was fined €20,000 for billing a consumer for third-party services she had disconnected: it violated her right of disconnection up to three times and obstructed the consumer inspection; it ended up accepting responsibility and paying €20,100 with the legal reduction. Movistar received a €20,000 fine after a debt collection agency harassed a consumer over a false debt of €1,741 — three lines she had never contracted. And Telefónica carries a €765,000 fine, confirmed by the Supreme Court, for limiting its

customers' right to receive paper bills. On the same list of grievances, consumer associations keep early-termination penalties under complaint as terms they consider abusive [2].

The common lesson is this book's: all these sanctions came because someone managed to document what was happening — and often after years of struggle. For every case that ends in a fine, how many die on the first call, when the customer cannot even prove they called? Consumer law recognizes the right to cancel in the same form the contract was made; what is missing is for the act of requesting it to leave a verifiable trail in the hands of the person requesting it.

The proof of existence applied

The same phone that performs the transaction can seal it [3].

1. The sealed cancellation

The customer performs the transaction and documents it in the moment: the recording of their own call — one's own conversation, lawful, as chapter 75 explains — the certified screenshot of the submitted cancellation form or of the chat with the bot, the confirmation email scanned [3]. Each piece receives its SHA-512 fingerprint on the device and its certain date. The cancellation request of the 3rd becomes a fact, whatever the company's system says on the 28th — and the PDF certificate with its QR code is attached to the consumer claim as is.

2. The contract and its versions

The terms published on the day you signed up, captured with t2p's built-in browser — the full image of the page and its source code, packaged and certified [3]; the contract scanned, up to 50 pages in a single PDF; and every unilateral modification, captured with its date. When the "forever" price rises in month eight, the comparison is arithmetic: this is what you said, this is what you say.

3. The call and the promise

The verbal promise — "we'll keep your price for another year if you stay" — recorded by the person receiving it, with the audio mode running in the background [3], and sealed on hanging up. The retention offer that later "is not on record" stops being unprovable. And the app's projects group the case's entire file — calls, captures, emails — ordered and dated [3].

4. The customer's file — and the symmetry

When the claim reaches the consumer authority or the court, the customer exports the .asice package: every piece with its fingerprint, its timestamp and the verification instructions, checkable by anyone without trusting anyone [3]. And the usual symmetry: the diligent company — the one that answers in two minutes, processes the cancellation on the 3rd and replies within five days — holds in its own anchored records the best defense against the unfounded claim and the best evidence of compliance with Law 10/2025 [1].

And the limits, for there are limits: the seal attests transactions — what you requested, when, what they answered — it does not decide the merits: whether the penalty was proper or the tariff correct is for the law and the judge to say. Recording covers only one's own conversations, never third parties', and the operator's data protection demands proportionality: you document the transaction, you do not hound the person answering. And no fingerprint replaces the formal claim within the deadline.

The promise and the peril

The promise: the customer exercising the rights of Law 10/2025 with the file in hand [1][3]; consumer claims resolved on dated facts rather than versions; bad practices — the returning charge, the cancellation that "is not on record" [2] — that stop being profitable because they leave a trail; and the honest company distinguishing itself with proof, not slogans.

The perils, with the same frankness.

The first is **the war of recordings**. If every call becomes an evidentiary minefield, the relationship degrades for everyone. The answer is proportionality: you document the transaction that matters — the

cancellation, the charge, the contractual promise — not every human interaction; and one's own recording is a shield for the dispute, never an instrument of harassment against the worker who answers.

The second is **the asymmetry that persists**. The customer with tools is still weaker than the company with a legal department. Proof levels the specific dispute; real protection is collective — consumer associations, authorities with sanctioning powers, and Law 10/2025 genuinely enforced [1][2].

The third is **ticket techno-solutionism**. No seal makes them pick up the phone in three minutes or makes cancellation simple: that is for the regulator to enforce. Proof makes the breach visible and claimable; giving it consequences is the job of whoever sanctions.

And the fourth is **the file's privacy**. The customer's file contains their consumer life: it must stay on their device — only fingerprints on the chain [3] — be disclosed selectively to the proper body, and never become a marketable profile.

It's time to act

Let the customer's word weigh as much as the company's database.

For you, the customer: document the transactions that matter — the cancellation, the unauthorized charge, the price promise — in the moment and with a certain date [3]; and claim within the deadline: Law 10/2025 gives you fifteen days' maximum response, and five for unauthorized charges [1].

For consumer associations: recommend the sealed file as the standard claiming practice [2][3]; every documented case is potential case law.

For the companies: anchor your service records — the Law 10/2025 tracking code [1] with a verifiable certain date is your proof of compliance and the end of the unfounded claim.

For consumer authorities: value the consumer's anchored evidence for what it is — a neutral certain date — and demand equally verifiable records from the companies [1][2].

For legislators: the rollout of Law 10/2025 [1] is the occasion: let every claim's identifying code carry a qualified timestamp, and let the cancellation record be verifiable by both parties.

The next time a voice says "call back tomorrow," let tomorrow have a certain date. For the customer who keeps their side and the company that serves, **time2act**.

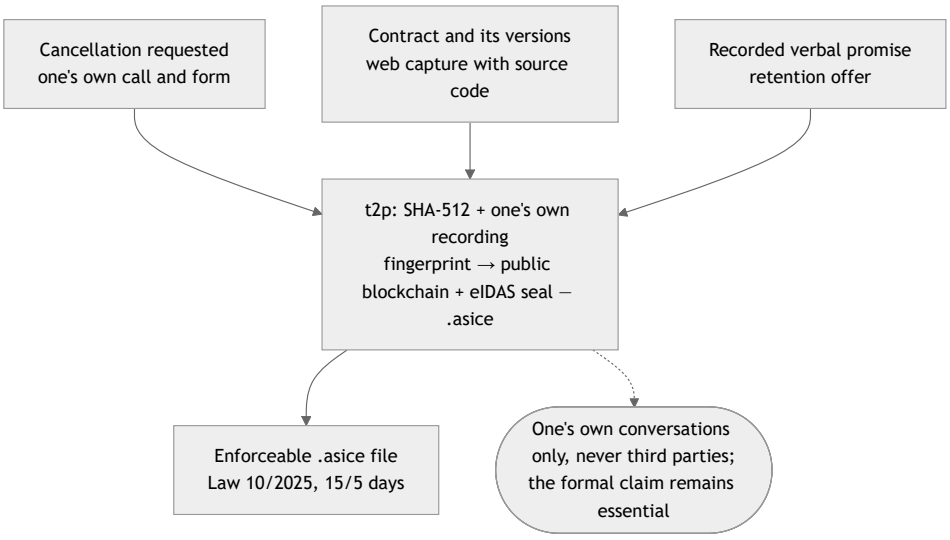


Figure 33.1 — The customer before the company with proof of existence (sealed with t2p — chapter 13): the sealed cancellation (recording of one's own call, certified capture of the form), the contract and its versions (web capture with source code), the recorded verbal promise and the .asice file exportable to the consumer authority. Law 10/2025 grants the rights (3 minutes, tracking code, 15/5 days, simple cancellation); the certain date makes them enforceable — and vouches for the diligent company. Limits: one's own conversations, proportionality with the operator, the formal claim remains indispensable. Sources: BOE (Law 10/2025); FACUA (Vodafone/Movistar/Telefónica fines).

[1] BOE — *Ley 10/2025, de 26 de diciembre, por la que se regulan los servicios de atención a la clientela* (BOE 27-12-2025; 12 months to adapt; 95% of calls in <3 minutes; identifying tracking code; response in 15 days — 5 for unauthorized charges; 15 days' notice before automatic renewals; simple cancellation with ID and contract number; ban on exclusively automated service); La Moncloa — *Ley de Servicios de Atención a la Clientela: objetivos y medidas*; Sepín — *Novedades de la Ley 10/2025*. https://www.boe.es/buscar/act.php?id=BOE-A-2025-26698 · https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/derechos-sociales-consumo-agenda-2030/paginas/2025/ley-servicios-atencion-al-cliente.aspx · https://blog.sepin.es/blog/2026/01/14/novedades-ley-10-2025-servicios-atencion-clientela/ ↩

[2] FACUA — *Multa de 20.000 euros a Vodafone: cobró a una socia de FACUA servicios de terceros que había desconectado* (right of disconnection violated up to three times; obstruction of the inspection; €20,100 with reduction); FACUA — *Multa a Movistar por acosar a través de una gestora de cobros a una socia por una deuda falsa* (€1,741 for lines never contracted); FACUA — *El Supremo confirma la multa de 765.000 € a Telefónica por limitar el derecho a recibir facturas en papel*; Xataka Móvil — *FACUA denuncia a Telefónica por la penalización por baja anticipada*. https://facua.org/noticias/multa-de-20-000-euros-a-vodafone-cobro-a-una-socia-de-facua-servicios-de-terceros-que-habia-desconectado/ · https://www.facua.org/es/noticia.php?Id=9279 · https://facua.org/noticias/el-supremo-confirma-la-multa-de-765000-de-la-junta-de-andalucia-a-telefonica-por-limitar-el-derecho-a-recibir-facturas-en-papel/ · https://www.xatakamovil.com/conectividad/facua-denuncia-a-telefonica-por-la-penalizacion-por-baja-anticipada-en-movistar-fusion ↩

[3] t2p — product manual (background audio recording; certified screen capture and recording with foreground app verification; web page capture with full image + source code; document scanner up to 50 pages; projects to group captures by case; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; the lawfulness of recording one's own conversations, in chapter 75; the cloned voice, in chapter 68. ↩

The Invoice and Late Payment: The Collection Clock

"Remember that time is money."

— Benjamin Franklin, "Advice to a Young Tradesman" (1748)

A carpentry workshop on an ordinary industrial estate. The order — custom furniture for a restaurant chain — rolls off the loading dock in the afternoon, the delivery note signed any old way on the van's hood. The invoice goes out by email the next day. And the wait begins: "next week," "it's pending validation," "we have no record of it." Meanwhile, the workshop's payroll does not wait: the cash box advances what the big client does not pay.

The scene has had a law for twenty years — and the law has data to be ashamed of [2]. Law 3/2004 sets maximum payment terms between companies: 60 calendar days at most by agreement, 30 by default from receipt of the goods or the service; it establishes automatic late-payment interest — no notice required, and in the first half of 2025 it stood at 11.15% — and a flat €40 in collection costs per invoice. The reality, according to the 2025 report of the Multisector Platform against Late Payment: 85% of large companies breach the legal deadlines, financing themselves partly at their suppliers' expense; public administrations pay on average at 70 days — more than double their 30-day limit — and the private sector at 67 [2].

Why is a law with automatic interest massively breached while almost no one claims? Because claiming demands proof: the date of delivery, the date the invoice was received, the exact start of the count. And those dates live in one-sided papers, disputed emails and delivery notes that get lost. This chapter deals with the evidentiary desert between issuance and collection — and how the certain date crosses it. International trade finance — letters of credit, bills of lading — is in chapter 54; accounting evidence for the auditor, in chapter 55; the pedigree of the goods themselves, in chapter 29; the private customer before the company, in the previous chapter. The key question: when, exactly, does the collection clock start? The answer: when the delivery and the invoice have a certain date opposable between private parties — the SHA-512 fingerprint on the device, the anchoring of the fingerprint to the public blockchain and the complementary RFC 3161/eIDAS timestamp of **t2p**, inside an .asice container [3]. And the state, without putting it in these words, has just proved this thesis right: it is called VeriFactu.

VeriFactu: the state builds half the proof

Royal Decree 1007/2023 — implementing the anti-fraud law — imposes a new architecture on all invoicing programs, the so-called SIF [1]. Every invoice must generate a standardized record with a **chained fingerprint**: the hash of each record incorporates the hash of the previous one, forming a chain in which any retrospective alteration sings out. Each record is electronically signed, and the invoice carries a **QR code** that allows it to be checked against the Tax Agency. There are two ways to comply: the "Verifactu" mode — immediate, voluntary submission of the records to the AEAT — or the unalterable preservation of the chain at the Inspectorate's disposal [1].

The calendar, after the one-year postponement approved by Royal Decree-Law 15/2025, of December 2: mandatory on January 1, 2027 for companies and July 1, 2027 for all other obligated parties. The penalties are not symbolic: up to €50,000 per fiscal year for whoever uses non-conforming software, and up to €150,000 for the manufacturers who sell it [1].

Pause a moment on what this means for this book's thesis: the Spanish state imposes by law chained hashes, unalterable records and verification by QR code. It is, piece by piece, the architecture of the proof of existence — applied to tax anti-fraud. It is no longer a theoretical proposal: it is the official gazette.

And now, the precision about what VeriFactu does not do — because this is where this chapter begins. First: it protects the tax office, not the creditor — the record attests that the invoice was issued and has not been altered, not that the client received it or that they owe you money. Second: it does not attest delivery — the delivery note, the receipt of the goods, the moment the payment term legally starts all remain outside the chain. And third: it is not a certain date designed to be opposable between private parties in a courtroom — it is evidence aimed at the tax administration. The missing layer — the one that looks to creditor and debtor, not to the taxpayer — is exactly what the neutral proof of existence adds [3].

The proof of existence applied

The same small business that issues the invoice can seal the whole journey to collection [3].

1. The delivery note with place and time

At the loading dock, the carrier or the technician photographs the signed delivery note with t2p: the forensic GPS requires a genuine fix — accuracy of about 50 meters, maximum age of about 30 seconds; a mocked location blocks certification outright [3] — and the SHA-512 fingerprint with its timestamp turns the delivery into a fact: what was delivered, where and when. If the industrial estate has no coverage, background certification resumes on its own when it returns, and the seal attests the moment of capture [3]. "The goods never arrived" runs into a record neither party can rewrite.

2. The invoice with a double certain date

The invoice, scanned and sealed at issuance — up to 50 pages in a single certified PDF [3] — and, above all, its dispatch captured: the email with the invoice attached, certified with screen and date [3]. Issuance and delivery, both with a certain date. The classic "the invoice never reached me" — the wild card that stops clocks — stops working: the invoice existed on such a day and was sent on such another, and the certificate with its QR code says so to anyone who cares to check.

3. The clock that starts by itself

With the delivery dated and the invoice delivered, the count under Law 3/2004 [2] is no longer arguable: 30 or 60 days from a certain fact, the automatic late-payment interest — 11.15% — and the €40 per invoice, computable to the cent from the exact day. The interest claim, which today almost no small business makes because it cannot fix day one, becomes arithmetic. And the negotiating position changes: the debtor knows the meter is running — and that it is on the record.

4. The demand and the payment order

When claiming becomes necessary, every step is dated: the payment demand sent and captured, the client's promise to pay saved with its date, and the entire file — delivery note, invoice, dispatch, demands — exported as an .asice package verifiable offline [3], ready for the payment-order procedure or for the fund that buys the receivables portfolio. And the usual symmetry: the client who paid on time holds in their own anchored records the proof against the duplicated claim or the miscalculated interest.

The limits, stated clearly: the seal attests dates and documents, not the quality of the work or the merits of a commercial dispute — if the goods were defective, the expert will say so. VeriFactu [1] and t2p [3] complement each other but must not be confused: one looks to the tax office, the other looks to the parties. And no fingerprint collects an invoice: it makes the claim enforceable — which is no small thing — but final collection depends on the court and the debtor's solvency.

The promise and the peril

The promise: small businesses claiming the interest the law has granted them for twenty years, because they can finally fix day one [2][3]; late payers paying sooner because the attested meter is running; payment-order proceedings resolved on verifiable files; and a country where VeriFactu [1] will have normalized an idea bigger than anti-fraud: that economic records carry chained hashes and a certain date — the cultural step toward the universal proof of existence.

The perils, with the same frankness.

The first is that **market power is not cured with hashes**. The small business often does not claim interest for fear of losing the big client, not for lack of proof. The certain date removes the evidentiary excuse, but the imbalance is economic: what is needed are late-payment sanction regimes and public observatories that make claiming unnecessary — so that paying late is expensive by default.

The second is **VeriFactu-solutionism**. Complying with the tax office is not paying on time: a company can have all the chained records in the world and pay at 120 days. Confusing tax anti-fraud with creditor protection would be misreading the reform [1].

The third is **the burden on the micro-business**. If documenting collection becomes added bureaucracy, the self-employed will not do it. The tool must be the ten-second gesture at the dock [3], not an administrative department — and the bulk of the effort, the fiscal chain, is already done by the invoicing software under legal obligation [1].

And the fourth is **the trail as a pressure weapon**. The same file that protects the creditor can serve to hound the good-faith debtor with industrialized claims. The counterweight is double: the proof's neutrality — it also attests who paid — and the judicial control of the payment-order procedure.

It's time to act

Let every invoice carry its clock, and every clock its day one.

For small businesses and the self-employed: seal the whole journey — the delivery note with forensic GPS, the sent invoice, the demands [3]; the day you claim, the clock of Law 3/2004 [2] will have an attested day one.

For accountants and advisors: make the anchored collection file part of the standard service; VeriFactu [1] already takes you to the fiscal hash chain — offer the client the between-parties layer.

For the big paying companies: anchor your receipts and payments — the proof that you pay on time is your best argument before the annual late-payment report [2].

For public administrations: paying at 70 days against a 30-day limit [2] is unsustainable; with receipt and payment records anchored and published by fingerprint, the average payment period will stop being a disputed figure.

For legislators: the late-payment sanction regime, pending for years, now has its evidentiary infrastructure ready [1][2][3]; and let the B2B electronic invoice now being rolled out incorporate the certain date of receipt, not just of issuance.

So that getting paid on time stops being a privilege, **time2act**.

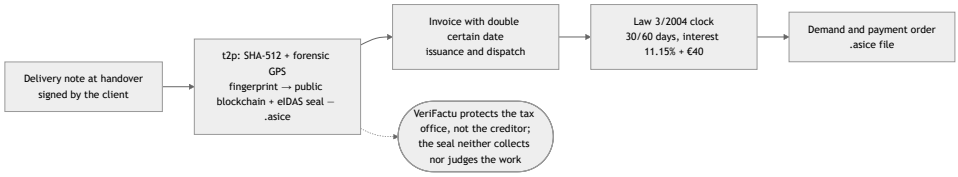


Figure 34.1 — The collection clock with proof of existence (sealed with t2p — chapter 13): the delivery note photographed with forensic GPS at delivery, the invoice with a double certain date (issuance and dispatch), the automatic count of Law 3/2004 (30/60 days; late-payment interest 11.15% + €40) and the .asice file for the payment-order procedure. Context: VeriFactu (RD 1007/2023, postponed to 2027 by RD-Law 15/2025) already imposes invoicing records with chained hashes and QR — the state builds half a proof of existence; t2p adds the layer opposable between private parties. Limits: the seal neither collects invoices nor judges the work; VeriFactu protects the tax office, not the creditor. Sources: AEAT; BOE; PMcM 2025 (85% of large companies in breach; public administrations at 70 days).

---

[1] AEAT — *Sistemas Informáticos de Facturación (SIF) y VERI*FACTU *(RD 1007/2023: standardized invoicing record with chained fingerprint — each record contains the previous one's hash — electronic signature and verification QR; "Verifactu" immediate voluntary submission mode and unalterable preservation; official FAQ); Guía Fiscal* — Verifactu 2026-2027 *and Davisa* — Guía Verifactu 2026* (calendar postponed by RD-Law 15/2025, of December 2: 1-1-2027 companies, 1-7-2027 the rest; penalties of up to €50,000/year for users and €150,000 for manufacturers). https://sede.agenciatributaria.gob.es/Sede/iva/sistemas-informaticos-facturacion-verifactu.html · https://sede.agenciatributaria.gob.es/Sede/iva/sistemas-informaticos-facturacion-verifactu/preguntas-frecuentes/sistemas-verifactu.html · https://guiafiscal.es/autonomos/verifactu-2026-guia/ · https://www.davisa.es/guia-verifactu-2026/ ↩

[2] BOE — *Ley 3/2004, de 29 de diciembre, de medidas de lucha contra la morosidad en las operaciones comerciales* (60 calendar days maximum by agreement between companies, 30 by default from receipt; automatic late-payment interest; €40 collection costs per invoice; public administrations 30 days — Law 15/2010); Togas — *Los intereses de demora en el primer semestre de 2025: 11,15 %* (Treasury Resolution of 30-12-2024); PMcM / Revista Limpiezas — *2025 late-payment report* (85% of large companies breach the deadlines; public administrations at 70 days on average, +3 in 2025; private sector at 67). https://www.boe.es/buscar/act.php?id=BOE-A-2004-21830 · https://togas.biz/articulos/articulo-profesionales--los-intereses-de-demora-en-las-operaciones-comerciales-en-el-primer-semestre-de-2025-quedan-fijados-en-un-11-15-por-resoluci-n-de-30-de-diciembre-de-2024-de-la-secretar-a-general-del-tesoro-y-financiacion-internacional/ · https://www.revistalimpiezas.es/actualidad/normativa-limpiezas/informe-morosidad-pmcm-el-85-de-las-grandes-companias-incumple-los-plazos-de-pago-y-se-financia-en-parte-a-costa-de-las-pymes-segun-la-pmcm_20260417.html · https://pmcm.es/informe-morosidad-estudio-plazos-pago-espana-2024/ ↩

[3] t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; background certification that resumes when coverage returns; document scanner up to 50 pages in one PDF; certified screen and email capture; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; international trade, in chapter 54; auditing, in chapter 55. ↩

Real Estate: The Condition of the Property and the Land Registry

"The first man who, having enclosed a piece of ground, bethought himself of saying 'This is mine,' and found people simple enough to believe him, was the real founder of civil society."

— Jean-Jacques Rousseau, Discourse on the Origin of Inequality (1755)

On an ordinary day, a homeowner discovers that his house has been put up for sale. He never sold it; someone has forged a deed — often a *quitclaim deed*, the American deed of renunciation — and filed it with the registry to transfer to themselves a property that is not theirs. The FBI's Boston office warned in April 2025 of the rise of this fraud, and the title insurance industry's association, the **American Land Title Association** (ALTA), documented that 28% of its companies had suffered **seller impersonation** attempts in 2023, with a favorite target: **vacant land** (85% of cases), because no one lives there to notice [5][6]. In many of those attempts, the fraudsters use fake notary credentials — 43% — or usurp real ones.

Modern real estate fraud, let us note, attacks not so much the chain of title as **the identity of the signer** and **the condition of the thing**. And here the system has a hole this chapter aims to fill. But it is worth beginning with a distinction the reader tends to blur and that runs through the whole chapter: one thing is the **legal title** — what creates or transfers ownership, the exclusive province of the registry and the notary — and quite another is the **evidence**: the proof that a document or a physical state existed, intact, at a given moment, in a given place, in the charge of a given person. The proof of existence does not replace the registry. It supplies the layer of evidence with a certain date that today is missing or contested.

Insecure property

The underlying problem is surprisingly massive. The World Bank, in work by its experts Laura Tuck and Wael Zakout, calculates that only about 30% of the world's population has its rights over land and housing legally registered; that is, close to 70% holds no documented title [1]. And the insecurity does not afflict poor countries alone. The Prindex initiative, the largest global survey on perceived security of tenure, found in 2024 that some 1.1 billion adults — 23%, up from 19% in 2020 — fear losing their land or home; in the United States this perceived insecurity has risen from 13% to 25% in four years, striking above all the young and renters [2]. A 2026 FAO report confirms it: only 35% of the world's land is formally documented [3].

Where there is opacity, there is fraud. The FBI recorded in 2025 more than 12,000 victims of real estate fraud in the United States, with losses of some $275 million, 59% more than the previous year [4]. Between 2019 and 2023 real estate fraud left 58,141 victims and $1.3 billion in losses [6]. And at the other end of the scale — that of the millions without title — insecurity is measured not in stolen dollars but in families who can be evicted from the land they work because they cannot prove their ownership.

Pioneers and failures: Georgia, Sweden, Honduras

Faced with this, some states have made a virtue of necessity. The pioneering case is the **Republic of Georgia**, a post-Soviet country scarred by registry corruption. In 2016, its National Agency of Public Registry, with the company Bitfury and with the economist **Hernando de Soto** — the great theorist of the dead capital of the untitled poor — as the project's godfather, began anchoring to the Bitcoin blockchain the fingerprint of every property certificate. By 2018, close to 1.5 million titles were sealed so that, in the agency's own words, "no one, not even the State, can alter their timestamp" [7]. In 2017 the pilot was extended to sales, mortgages and rentals [8]. The paradox is beautiful: a state that gives its citizens a proof it can itself no longer falsify.

Sweden went down a similar road: its registry, the Lantmäteriet, with the company ChromaWay, calculated that sealing real estate transactions could save taxpayers more than **€100 million a year** and cut the time between signing and transfer from three-to-six months to a few days [9]. But not all the stories end well, and there is a lesson here. Honduras announced in 2015, with the company Factom, the first "tamper-proof" national land registry — motivated, precisely, by officials hacking the state database to steal properties for themselves. By the end of that year, the project had run aground on "political questions" and was never deployed [10]. The technology promised to armor property; the lack of political

will left it a letter of intent. It is proof that the seal is a necessary condition, but not a sufficient one: the registry, the notary or the court must accept it.

The proof of existence applied

Let us see how **t2p** fills the evidence gap in three very different scenarios.

The first, everyday and universal: **rental deposit disputes**. The classic "it was already broken / no, you broke it" is one of the most frequent causes of litigation between landlord and tenant, to the point that seventeen US states require by law a condition report at move-in before a deposit can be legally withheld. With **t2p**, the tenant or the manager photographs the state of the flat and seals it on the spot: a **SHA-512 fingerprint** on the device, **forensic GPS** proving which flat and when, an **eIDAS timestamp**, a **proof of life** tying the capture to a real person present — goodbye to recycled old photos — and everything packaged into an **.asice container** that the other party or a judge can verify offline.

The second: **deeds, contracts and due diligence**. Sealing the fingerprint of a draft deed, an earnest-money contract or an appraisal report gives them a **certain date**, making backdating impossible. In the seller-impersonation vector, capturing the moment of identification with proof of life and GPS adds a trail that forged notary credentials cannot fabricate after the fact.

The third, the most transformative: **weak-registry contexts**, as we have seen. Where no reliable registry exists, sealing georeferenced photographs of occupation, improvements and boundaries creates a **verifiable history of possession**. It is not a title — only the state can create that — but it is evidence with a certain date, opposable, in the spirit of what Georgia did. It is, for millions of families, the difference between a mute possession and a possession that can be proven.

The limits must be stated clearly, faithful to the duality of "the promise and the peril." "Garbage in, garbage on-chain": if the photo lies or the source document is false, the chain only certifies that that lie existed intact at that instant. The proof of existence demonstrates existence, integrity, time, place and person; it does **not** demonstrate the prior physical truth — that the wall was sound before, or that the seller is the real owner. And, as in Honduras, its full value appears only when the registry, the notary or the court accepts this evidence.

The promise and the peril

The promise is double. For rich countries, closing the door on deed fraud and deposit disputes, and giving a certain date to every real estate document. For countries with weak registries, offering hundreds of millions of people a way to attest their possession without waiting for the state to build a perfect cadastre. In both cases, more legal certainty at less cost and with fewer intermediaries.

The peril demands nuance. First, the dangerous confusion between evidence and title: selling the proof of existence as if it transferred ownership would be a fraud in itself. Second, privacy: geolocated photos of a home's interior and liveness biometrics are sensitive data demanding scrupulous design. And third, the lesson of Honduras: without political will and institutional recognition, the best technology remains a promise. The proof of existence is a necessary layer within an architecture that always includes the registry and the law.

It's time to act

For **tenants and landlords**: seal the condition of the dwelling at move-in and move-out, so a deposit dispute is not settled by the loudest voice. For **buyers and title insurers**: demand evidence with a certain date and the seller's proof of life, cutting off impersonation. For **registries and notaries**: recognize the seal's evidentiary value and integrate it, as Georgia did. For **governments with weak registries**: offer a layer of possession evidence while the cadastre is built, without making the population wait. For **regulators**: protect simultaneously the integrity of the proof and the privacy of homes.

A homeowner can discover one day that his house has been "sold" without him. Millions of families live their whole lives unable to prove that the land they work is theirs. Between verifiable property and property resting on trust in a manipulable document, the frontier of security is increasingly drawn. Giving

everyone the proof of what is theirs — or of the condition in which they received it — is, here too, **time2act**.

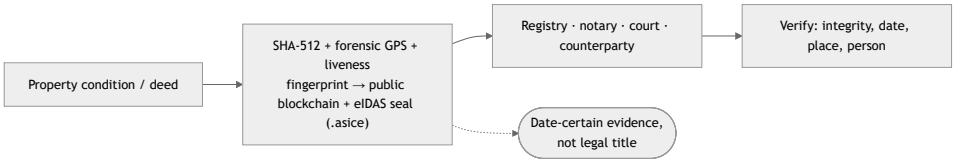


Figure 35.1 — The condition of the dwelling (rental) or a document (deed, contract) is captured and sealed on the device with SHA-512, forensic GPS, an eIDAS timestamp and proof of life, anchored to the public blockchain and packaged into an .asice container verifiable offline, ready for the registry, the notary, the court or the other party. The proof of existence supplies evidence with a certain date, not the legal title (which the registry creates). Context bar: ~70% of the world's population without registered title; ~1.1 billion people insecure about their property; Georgia, ~1.5 million titles anchored. Sources: World Bank; Prindex; Exonum/Bitfury.

---

[1] World Bank — Laura Tuck and Wael Zakout, *7 reasons for land and property rights to be at the top of the global agenda* (~30% of the population with registered title). https://blogs.worldbank.org/en/voices/7-reasons-land-and-property-rights-be-top-global-agenda ↩

[2] Prindex — *Tenure insecurity on the rise* (2024; ~1.1 billion adults, 23%, insecure about their land/housing). https://www.prindex.net/news-and-stories/tenure-insecurity-on-the-rise/ ↩

[3] FAO (SOFA report, via UN News) — *Global report highlights slow progress in expanding secure land tenure* (~35% of land documented; >1,100 M fear losing it). https://news.un.org/en/story/2026/02/1167037 ↩

[4] FBI IC3 (via the National Association of Realtors) — *Online real estate fraud climbed to $275M in 2025* (12,368 victims). https://www.nar.realtor/magazine/real-estate-news/online-real-estate-fraud-climbed-to-275m-in-2025-fbi-says ↩

[5] American Land Title Association (ALTA) — *Seller Impersonation Fraud Attempts Increase, Study Shows* (July 2024; 28%; vacant land 85%; fake notary credentials 43%). https://www.alta.org/news-and-publications/news/20240730-Seller-Impersonation-Fraud-Attempts-Increase-Study-Shows ↩

[6] FBI Boston (via ALTA) — *FBI Boston Issues Quit Claim Deed Fraud Warning* (April 2025; real estate fraud 2019-2023: 58,141 victims / $1.3 billion). https://www.alta.org/news-and-publications/news/20250410-FBI-Boston-Issues-Quit-Claim-Deed-Fraud-Warning ↩

[7] Exonum / Bitfury — *The story of Georgia* (NAPR; ~1.5 million titles anchored to Bitcoin; "no one can alter the timestamp"). https://exonum.com/story-georgia ↩

[8] Eurasianet — *Georgia authorities use blockchain technology for developing land registry* (extension Feb-2017; Hernando de Soto, adviser). https://eurasianet.org/georgia-authorities-use-blockchain-technology-for-developing-land-registry ↩

[9] Computer Weekly — *Sweden trials blockchain for land registry management* (Lantmäteriet + ChromaWay; >€100 million/year in savings). https://www.computerweekly.com/news/450421958/Sweden-trials-blockchain-for-land-registry-management ↩

[10] CoinDesk — *Blockchain land title project stalls in Honduras* (Factom; stalled on political questions, 2015). https://www.coindesk.com/markets/2015/12/26/blockchain-land-title-project-stalls-in-honduras ↩

Renting and Housing: The State of Things, With a Certain Date

"The house is our corner of the world. As has often been said, it is our first universe."

— Gaston Bachelard, The Poetics of Space (1957)

You are on a well-known property portal and the cursor stops over a listing that looks like a gift: a pretty flat, well located, clearly below the neighborhood's price. The "owner" writes with urgency — there are other interested parties, a deposit is needed to reserve it — and there is no way to visit first: he is away for work. Just an account number to transfer the money. That is how it began, for a victim flat-hunting in the center of Valladolid, the story the Guardia Civil ended up christening Operation ISLEDA (2026): a "nationwide real estate scam" ring, with fake listings scattered across the country and victims all over Spain; three arrests, aged 19 to 34, in San Fernando (Cádiz); and one detail that captures the age — the receiving bank accounts had been opened with personal documents stolen in earlier fake job-offer scams. The charges: fraud, money laundering and criminal organization [1]. The pattern is always the same: bait price, an owner who never shows his face and haste to make you pay before you can check anything [2].

Chapter 35 dealt with ownership: the title, the registry, the land. This chapter deals with what happens inside: the life of the rental contract, where almost everything is argued after the fact. Whether the flat existed. What condition it was in the day you got the keys. When you reported the damp. When the squatter got in. The key question: how do you prove the state of things when each side remembers a different version? The answer cannot be memory: it is the certain date in the moment — **t2p**: the SHA-512 fingerprint of every photo, inventory record and message, computed on the device, with an RFC 3161/eIDAS timestamp and anchoring to the public blockchain inside an .asice container [5]. And with total neutrality: the same tool shelters the tenant, the landlord and anyone acting in good faith.

The flat that does not exist

The mechanics of the fake-listing scam are well documented: nonexistent flats — or flats that do not belong to the advertiser — posted with attractive photos and below-market prices to build a queue of victims; an "owner" who often claims to be abroad; and fast-paced communication pressing for payment of a reservation or the first month before any visit [1][2]. Operation ISLEDA, which came to light in the spring of 2026, showed its industrial maturity: victims across the territory, listings distributed nationwide, stolen identities used to open the receiving accounts and a criminal-organization structure [1].

The underlying asymmetry explains why it works. To set up the scam all you need are pretty pictures, persuasive copy and an email address; the cost is trivial next to the takings. By contrast, to prove you are the genuine owner there is today no simple gesture: the tenant cannot tell the legitimate advertiser from the impostor until it is too late, and the honest landlord competes against phantom flats cheaper than his own.

The structural way out is the verifiable listing. The real owner can seal the offer — linked to his accreditation of title, the registry chain of chapter 35 — with a certain date [5]; the portal can display the accredited-listing badge; and the tenant learns the inverse rule: an offer without a seal and with haste, suspect by default [2]. The scammer can copy the photos, but he cannot attest either the title or the continuity over time. Without the badge, his listing stops competing; with the rule learned, he stops collecting.

The deposit war

When the contract ends, the other battle begins. The framework is clear on paper: the deposit must be returned within one month of the contract's end, and after that deadline it accrues statutory interest [3]; the valid grounds for withholding are enumerated — rent or utility debts, and damage exceeding normal wear and tear [3]. Paint worn by time justifies withholding nothing; a broken tile or a forced door does.

The key, however, is procedural: the courts tend to rule for the tenant when the landlord provides insufficient proof — invoices, damage reports — to justify the withholding; and for the landlord when he substantiates real expenses from misuse [3]. Translated: the deposit is won not by whoever is right, but by whoever has proof.

And here appears the missing ritual. The experts recommend formalizing a signed settlement agreement, with a detailed check of the inventory at the handover of keys [3]; in practice it is done rarely, late and without a reliable date — loose photos on each person's phone, scattered messages, clashing memories.

The sealed version of the ritual fixes it at the root: on key day, both parties photograph the dwelling — every room, every pre-existing defect — and seal the images and the inventory record on the spot, from both devices: SHA-512 fingerprint, certain date, anchoring [5]. On move-out day, the same walk-through. The argument over "how the flat was" disappears: what remains is the comparison between two series of images with attested integrity. The tenant cannot deny the hole he made, nor can the landlord pin on him the crack that was already there.

The 48 hours and the 15 days

Squatting is the market's most poisoned conflict, and its new legal framework has made time the central element. Organic Law 1/2025, on the efficiency of the public justice service, in force since April 3, 2025, routes breaking-and-entering of a dwelling (Article 202 of the Criminal Code: occupying someone's habitual residence) and usurpation of property (Article 245: occupying a property that is not a habitual residence) through fast-track trials, with resolution times on the order of fifteen days; if the occupation is detected within the first 48 hours, the security forces can evict without waiting for a court order, as a flagrant offense; and on this track the prior vulnerability report is abolished [4].

Notice what the entire fast track hangs on: dates. When did the squatter get in? When did the owner discover it? Was the property a habitual residence or a second home? When was it reported? Each answer moves the case from one article to another, from one track to another, from immediate eviction to the long procedure. And today these questions are answered with testimony and circumstantial signs — the most fragile evidentiary material there is.

With proof of existence, they are answered with chronology. The owner of an empty dwelling can establish the baseline with periodic condition checks — the monthly visit, the photo with a certain date [5]; if one day he finds the door forced, the discovery is dated to the minute and the report carries an anchored time: the 48-hour flagrancy and the fifteen-day trial [4] operate on dated facts, not estimates.

And the symmetry is indispensable, and must be said with the same emphasis: the same tool protects the legitimate tenant from being treated as a squatter. The contract, the rent receipts and the utility payments sealed on his device are his chronology of lawful possession [5], the shield against abusive eviction or unfounded accusation. The certain date takes no side: it distinguishes illegal occupation from legal possession, whichever side has to prove it.

Damp, breakdowns and the clock

The third war of renting is neither the move-in nor the move-out: it is the meanwhile. The damp stain that grows, the boiler that fails in winter, the repair that "is coming." Legally, these disputes are pure chronologies: when the landlord was notified, through which channel, what he answered, how long the repair took, how the damage evolved in the meantime. Whoever cannot date the notice cannot demand; whoever cannot date the response cannot defend themselves.

Imagine the sequence with proof of existence. The tenant spots the damp in November and sends the sealed notice — fingerprint and certain date of the message [5]. The landlord's reply, also sealed, arrives the next day: inspection next week. The technician documents the damage with dated photos; the repair goes on record; and if the stain returns in January, the anchored photo series shows the whole progression. Neither party depends on memory anymore: the tenant who gave timely notice can demand; the landlord who repaired diligently can prove it; and whoever stayed silent — whichever side it was — is exposed by their own silence, because the gaps in the chronology are dated too [5].

In a contract of successive performance, justice is a matter of the clock. And the clock, to be fair, must be auditable by anyone: that is exactly what the seal provides.

The proof of existence applied

An ordinary handover of keys: two phones photographing the same rooms, ten minutes of ritual — and a year of dispute saved.

1. The sealed inventory

On move-in day, tenant and landlord walk the flat together: photos of every room and of the pre-existing defects, a brief inventory record, and everything sealed on the spot from both parties' devices — SHA-512 fingerprint, certain date, anchoring to the public blockchain [5]. On move-out day, the same walk-through. The deposit is settled on the comparison of two image series with attested integrity, not on memories: the courts already reward whoever brings proof [3]; the seal ensures both parties have it, of equal quality, from day one.

2. The verifiable listing

The legitimate owner seals the offer with his accreditation of title (chapter 35) and the portal distinguishes it with the verified-listing badge [5]. Faced with the "too cheap" flat, the owner abroad and the rush [1][2], the tenant applies a three-word rule: no seal, no money. Operation ISLEDA took a complaint, an investigation and months of police work to dismantle [1]; a market of accredited listings would have left it without victims from the first click.

3. The chronology of possession

For the owner of an empty dwelling: periodic, sealed condition checks that establish the baseline; if one day he finds the door forced, the discovery is anchored to the minute, the report as well, and the fast track of LO 1/2025 — the 48 hours, the trial in about fifteen days [4] — runs on an indisputable chronology [5]. For the legitimate tenant: the sealed contract, receipts and utilities are the shield that prevents anyone from painting him as a squatter. The same fingerprint serves both faces of the conflict: that is precisely its virtue.

4. The dialogue with a certain date

The breakdown notice sealed, the landlord's response sealed, the damage's progression in an anchored photo series [5]: the damp dispute stops being an "I told you" against a "you never told me." Whoever gave timely notice can demand; whoever repaired diligently can prove it; and the silence of whoever did neither is dated with the same precision as everything else.

What the seal does not do must also be said: it attests the condition and the moment, but it does not apportion blame or interpret the Urban Leases Act — that is for judges; a sealed photo does not replace physical inspection when an expert is needed; and images of a home's interior are intimate data — only the fingerprint is anchored, the content stays on the parties' devices and is revealed only in case of dispute [5].

The promise and the peril

The promise is a rental market where the fake listing finds no victims [1][2], the deposit is settled in ten minutes of comparison rather than months of war [3], squatting and legitimate possession are distinguished by chronology [4], and repairs run on a clock both parties can audit. Fewer courts, less fear, more trust between strangers — in the most everyday contract there is.

The perils demand the same honesty.

The first is **the imbalance of power**. If only one side seals, the tool becomes a weapon. A market where large holders document industrially and the vulnerable tenant does not aggravates the asymmetry that already exists. That is why the ritual must be bilateral by design — the inventory sealed from both devices — and the tool free or trivially cheap for the weaker party.

The second is **the privacy of the home**. The photos of your house are the map of your intimacy. No image should leave the device: only fingerprints on the chain, disclosure only upon dispute, revocability when the contract ends and no central registry of interiors [5].

The third is **the bureaucratic barrier**. If the seal becomes one more requirement for getting a flat — a "documentation score" that expels those who have least — we will have worsened the problem we wanted to solve. The seal must settle disputes, not filter people: simple, free and optional for the renter.

And the fourth, said without hedging: **this is not housing policy**. No fingerprint lowers the rent or builds a single flat. The housing crisis is one of prices, supply and public policy; this chapter only clears away the layer of conflicts that makes it crueler still. Confusing the two would be selling smoke over people's roofs.

It's time to act

In the most important contract of daily life, memory cannot be the only proof.

For tenants: seal the move-in inventory even if the landlord does not; keep the contract, the receipts and every breakdown notice with a certain date [3][5]; and before a listing with haste and a bait price, no seal and no visit — not one euro [1][2].

For landlords: a bilateral sealed inventory on key day [3]; for empty flats, the periodic condition check [4][5]; and accredit your listing — the badge sets you apart from the scammer copying your photos [1].

For property portals: adopt the verified-listing badge with a title seal and a certain date [1][2][5]; the cost is minimal and the whole market gains credibility.

For property managers and landlords' chambers: standardize the sealed inventory record as a basic service of professionalized renting [3].

For judges and police forces: weigh the sealed chronology on the fast track of LO 1/2025 [4] — the 48 hours and the fifteen days need reliable dates more than anything else.

For legislators: grant a presumption of authenticity to the bilateral sealed inventory, and see to it that no evidentiary tool ever becomes a barrier to housing.

The Valladolid victim lost her money over a flat that did not exist [1]. May the next fake listing find no one to deceive, and may the next deposit be returned without a war. For a market where the state of things carries a certain date, **time2act**.

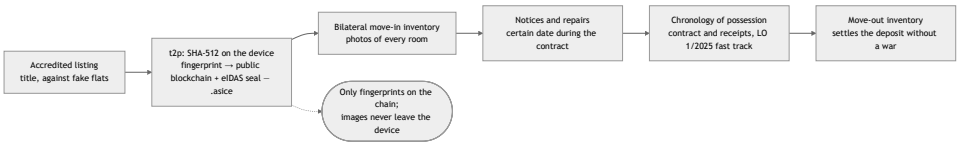


Figure 36.1 — The rental cycle with proof of existence: the accredited listing (title + seal, against the nonexistent-flat scam), the move-in inventory sealed bilaterally, the notices and repairs with a certain date during the contract, the chronology of possession (the tenant's contract and receipts; the owner's periodic condition checks; LO 1/2025 fast track) and the move-out inventory that settles the deposit without a war. Only the SHA-512 fingerprints are anchored to the public blockchain (RFC 3161/eIDAS, .asice); the images of the home never leave the devices. Sources: Guardia Civil (Op. ISLEDA); Idealista/Maldita; LO 1/2025.

---

[1] Libertad Digital — *Cae en Cádiz una red criminal por una «estafa inmobiliaria a nivel nacional»* (May 6, 2026; Operation ISLEDA; 3 arrests aged 19-34 in San Fernando; uncovered in Valladolid; accounts opened with identities stolen in fake job offers; fraud, laundering and criminal organization); Cádiz Directo — *La Guardia Civil destapa en San Fernando una trama nacional de falsas viviendas de alquiler*. https://www.libertaddigital.com/andalucia/2026-05-06/san-fernando-cadiz-valladolid-guardia-civil-detiene-tres-hombres-acusados-de-estafa-inmobiliaria-conocido-portal-web-alquiler-y-venta-de-pisos-utilizaban-sus-datos-para-cuentas-falsas-victimas-por-todo-el-pais-destapado-en-valladolid-1b-7400307/ · https://www.cadizdirecto.com/cadiz/la-guardia-civil-destapa-en-san-fernando-una-trama-nacional-de-falsas-viviendas-de-alquiler ↩

[2] Idealista/News — *La última estafa en los anuncios de alquiler y los consejos para no caer en ella* (bait prices, owners "abroad," payments before the visit); Maldita.es — *Anuncios de alquiler falsos e inmobiliarias fantasma: cómo reconocer timos inmobiliarios*. https://www.idealista.com/news/inmobiliario/vivienda/2022/11/18/800160-la-ultima-estafa-en-los-anuncios-de-alquiler-y-los-consejos-para-no-caer-en-ella · https://maldita.es/malditobulo/20230830/inmobiliaria-fantasma-alquiler-piso-fraude-verano/ ↩

[3] Idealista/News — *Todo lo que tienes que saber sobre la devolución de la fianza del alquiler* (one-month deadline; statutory interest; settlement agreement with inventory); Ático Jurídico — *Devolución de la fianza del alquiler* (courts demand proof from whoever withholds: invoices, reports); Fotocasa Life — *Cómo proceder si no nos devuelven la fianza del alquiler*. https://www.idealista.com/news/inmobiliario/vivienda/2022/07/03/797535-todo-lo-que-tienes-que-saber-sobre-la-devolucion-de-la-fianza-del-alquiler · https://aticojuridico.com/devolucion-fianza-alquiler/ · https://www.fotocasa.es/fotocasa-life/alquiler/como-proceder-si-no-nos-devuelven-la-fianza/ ↩

[4] Housage — *Ley antiokupas 2025: desalojo en 15 días, cómo funciona y a quién no protege* (LO 1/2025 in force April 3, 2025; breaking-and-entering Art. 202 CP and usurpation Art. 245 CP via fast-track trial; police eviction within the first 48 hours;

abolition of the vulnerability report on this track); Blegal — *La Ley 1/2025 agiliza el desalojo de okupas*; Vento Abogados — *Juicio rápido por okupación: la nueva regulación de la LO 1/2025*. https://housage.es/articulos/actualidad-inmobiliaria/ley-antiokupas-2025-desalojo-en-15-dias-como-funciona-y-a-quien-no-protege/ · https://www.blegal.eu/la-reforma-de-la-ley-organica-1-2025-agilizara-los-juicios-rapidos-para-facilitar-el-desalojo-de-okupas/ · https://vento.es/juicio-rapido-por-okupacion/ ↩

[5] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; content never leaves the device; voluntary disclosure; the gaps in the series are dated too). See chapters 13 and 15; for title and the registry, chapter 35. ↩

The Homeowners' Association: Minutes, Special Levies and Agreements With a Certain Date

"Cities have the capability of providing something for everybody, only because, and only when, they are created by everybody."

— Jane Jacobs, The Death and Life of Great American Cities (1961)

The annual meeting, in the lobby or the room under the stairs: folding chairs, the minute book under the property manager's arm, the eternal argument over the special levy for the elevator. Hands go up, the count is done in a hurry, someone protests and someone else is already on the next item. Three weeks later the minutes arrive — and more than one of those present does not recognize what they say was voted.

Millions of people live under condominium rules, and their entire collective life — works, money, rules — hangs on the flimsiest documentary chain: a notice that may or may not have arrived, a vote nobody recorded, minutes written up afterward by a single person, a minute book that lives in a drawer. When there is trust, everything works; when it breaks — and in a residents' community it breaks often — everything turns into "your word against the president's."

This chapter is about the country's smallest democracy and its memory. The title deed and the registry are in chapter 35; the life of the rental contract, in chapter 36; the relationship with public administration, in chapter 44; and if a parking space ever has to be assigned by lottery, chapter 45 explains how to do it irrefutably. The key question: how do you govern a building when the only official memory is drafted and kept by a single hand? The answer: giving a certain date to every piece — the notice, the minutes, the accounts, the communication — with **t2p**: hash on the device, timestamp and public anchor — the tripod of chapter 13 [3]. Within reach of any resident — and of the honest property manager, who is the first to benefit.

The minute book in a single hand

The real case that illustrates it [1]: a property manager in Gijón, sentenced to a year and a half in prison — conviction upheld — for appropriating €2,984.04 from the community's account. The method: deceptive maneuvers in the minute book, getting former presidents to sign documents that supposedly settled formalities with the bank; to cover his tracks, he produced receipts for works and locksmithing that corresponded to no community expense. And the detail that matters to this book: the fraud was only discovered when the president and treasurer changed, and the new team went through the accounts minutely.

The crime is the exception — the vast majority of property managers do a thankless job well — but the architecture that allowed it is the norm: the minute book as the only official memory, drafted and kept by one party; accounts that are only examined when someone gets suspicious; signatures requested "for a formality." The Criminal Code gives it a name once everything has happened — the misappropriation of Article 253 [1]; what does not exist is anything to make it visible beforehand.

And the civil conflict is far more frequent than the criminal one. The Horizontal Property Act allows agreements that are unlawful, harmful or abusive to be challenged [2] — but challenging demands proving what was convened, what was said and what exactly was voted. Without a record, the challenge is an act of faith; and the defense of a board that did everything right is equally at the mercy of memories.

The proof of existence applied

The same meeting, the same manager — and every piece of community life with a certain date [3].

1. The notice that is on record

The meeting notice with the agenda, scanned and sealed when sent — and the sending itself captured with a certain date: the email, the noticeboard or the mailbox photographed with forensic GPS [3]. "It never reached me" and "it was convened out of time" stop being debatable: the notice existed on such a day, said this and was communicated through this channel. It protects the resident — and protects the board that convenes properly.

2. The same-day minutes

The vote, counted aloud and — if the meeting so agrees — recorded [3]; the minutes, drafted and sealed the same day or the next, not three weeks later. Each set of minutes receives its SHA-512 fingerprint with a timestamp, and the minute book becomes a chain of anchored versions that no one — manager, president or resident — can rewrite retroactively [3]. The challenge under the Horizontal Property Act [2] moves on to debating the substance, not the facts; and the maneuver of the Gijón case [1] — touching the book after the fact — becomes impossible to perform in silence.

3. The traceable special levy

The works file, grouped into a project in the app [3]: the quotes received — scanned, up to 50 pages each — the approval agreement with its sealed minutes, the invoices, the certifications. Every euro of the levy, from vote to payment, with a certain date. The annual accounts, checkable by any resident without asking permission — and the honest manager, armored: his anchored records are his permanent audit.

4. Reliable neighborhood communication

The communications that matter — the demand to the manager, the notice of the defect that will later become an insurance claim, the formal complaint — sent and captured with a certain date [3], without a burofax or a notary for every trifle. And when the conflict escalates, the entire file is exported as an .asice package verifiable offline [3], ready for the lawyer or the court.

The limits, so as not to sell smoke: the seal attests notices, minutes and accounts — it does not resolve the substance of any neighborly dispute: whether the levy was fair or the works necessary is decided by the majorities and, if need be, the judge [2]. Recording meetings requires the meeting's own agreement and proportionality: you document the collective decision, never the residents' lives. And minutes contain personal data: publication is by fingerprint; the content stays within the community.

The promise and the peril

The promise: communities where trust does not depend on anyone's memory [1][3]; challenges that debate the substance on certain facts [2]; honest managers holding the profession's best credential — verifiable records; and the country's smallest democracy finally running with the traceability we demand of the big ones.

Now for the perils.

The first is **the surveilled community**. If every stairwell becomes a field of crossed recordings, more will have been lost than gained. The line is sharp: the acts of governance are documented — notice, vote, minutes, accounts — by collective agreement; never the coexistence, never the neighbors among themselves.

The second is **the war of blockades**. The litigious neighbor with tools can flood the community with documented challenges. The counterweight is the same symmetry: the board that does things right holds the file that proves it, and the unfounded challenge dies of facts.

The third is **the digital divide on the landing**. The community belongs to everyone, including those without a smartphone. No system may exclude: paper still counts — it is scanned and anchored — and the certifier is the community as a body, not each resident on their own.

The fourth peril lives on the landing itself: expecting technology to play the good neighbor. No fingerprint fixes a poisoned community or pays the levy of those who cannot. The certain date pacifies the layer of facts; coexistence and solidarity among neighbors are — and will remain — human work.

It's time to act

Ten minutes of agreement at the first meeting save years of arguments.

For presidents and boards: agree on the protocol — notices, minutes and accounts sealed [3]; it is one short vote that changes the community's life.

For property managers: anchor the minute book and the accounts you manage [1][3] — your work well done deserves to be beyond dispute, and the whole profession gains when fraud has nowhere to hide.

For residents: keep with a certain date what affects you — the notice, the defect report, the communication to the manager [3]; on the day of the challenge [2], facts will be worth more than memories.

For the associations of property managers: make the anchored record a professional standard and a mark of quality — the Gijón case [1] is the counterexample the profession cannot afford.

For legislators: the Horizontal Property Act already admits remote meetings [2]; the natural next step is the electronic minute book with a verifiable certain date — the public faith of community life, at a cost of cents.

At the next meeting, when the hands go up, let the count have a certain date — and let next year's argument be about the elevator, not about what was voted. For the democracy of the landing, **time2act**.

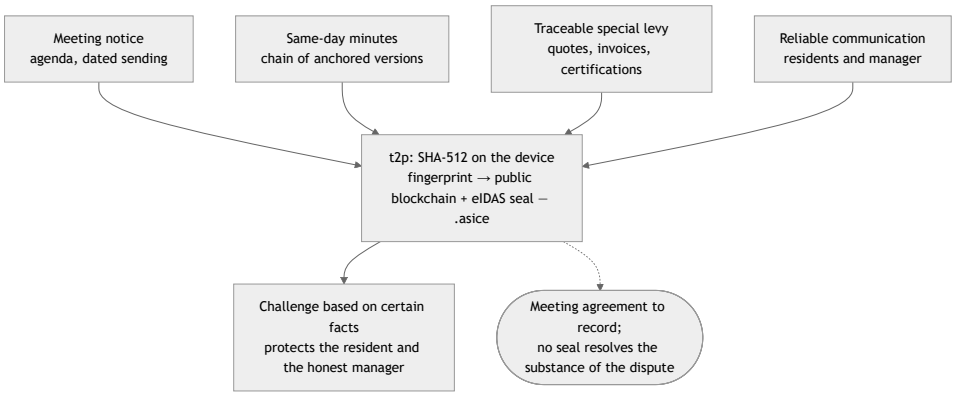


Figure 37.1 — The homeowners' association with proof of existence (sealed with t2p — chapter 13): the notice with a certain sending date, the minutes sealed the same day (the minute book as a chain of anchored versions), the traceable levy (quotes, agreement, invoices, certifications in one project) and reliable communication between residents and manager. Symmetry: it protects the challenging resident and the honest manager. Limits: meeting agreements to record, publication only by fingerprint (personal data), and no seal resolves the substance of disputes (Horizontal Property Act). Sources: Gijón conviction (Art. 253 CP); Horizontal Property Act.

---

[1] Vecindia — *Condenado administrador de fincas por apropiación indebida* (Gijón: 1 year and a half in prison upheld; €2,984.04 from the community account; deceptive maneuvers in the minute book getting former presidents to sign; fake works and locksmith receipts as cover; discovered when the president and treasurer changed); ADADE — *El Tribunal Supremo ratifica la condena por apropiación indebida a un administrador de fincas*; El Derecho — *Responsabilidad penal del administrador de fincas* (Art. 253 CP). https://vecindia.es/comunidades-de-propietarios/condenado-administrador-fincas-apropiacion-indebida/ · https://www.adade.es/comunicacion/sala-de-prensa/el-tribunal-supremo-ratifica-la-condena-por-apropiacion-indebida-a-un-administrador-de-fincas · https://elderecho.com/cuando-puede-concurrir-responsabilidad-penal-en-el-administrador-de-fincas-por-su-actuacion-en-una-comunidad-de-propietarios/ ↩

[2] Horizontal Property Act (challenge of agreements that are unlawful, harmful or abusive; notices and minutes; remote meetings enabled by recent reform); Mundojurídico — *Responsabilidad del administrador de la Comunidad*; Oliva-Ayala — *Falsedad documental y certificaciones de juntas*. https://mundojuridico.info/responsabilidad-del-administrador-de-la-comunidad/ · https://www.oliva-ayala.es/wp-content/uploads/2024/03/Alcacer-2024-LLPenal-Falsedades.pdf ↩

[3] t2p — product manual (document scanner up to 50 pages in one certified PDF; audio recording; certified screen and email capture; photo with forensic GPS; projects to group a works file; certify any file — only the fingerprint travels; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; property and the registry, in chapter 35; renting, in chapter 36. ↩

Home Renovations: The Advance, the Works and the Proof

"When we build, let us think that we build for ever."

— John Ruskin, The Seven Lamps of Architecture (1849)

The kitchen, half done: countertops ripped out, pipes exposed, a layer of dust on every surface. In the middle of the mess, a phone that no longer answers. The last message said: "the advance is paid, the materials are ordered — I'll be back tomorrow." The return date has been repeating ever since: always tomorrow.

It is no exception: it is a documented script [1]. The Guardia Civil and the police periodically investigate the same pattern: the attractive ad on a digital portal, the suspiciously cheap quote, the first payment up front "for the purchase of materials," and the disappearance. In Vigo and its district, a single investigation counted at least six victims with losses of up to €10,000 each; in other cases, accumulated losses have exceeded €77,000; and in Guadalajara, a resident lost his savings to a fake bricklayer from the internet until the Guardia Civil cornered the scammer in Madrid [1].

And the nuance that motivates the entire chapter: for every scammer there are hundreds of honest professionals who suffer the same evidentiary void from the other side — the client who haggles over the final invoice alleging defects, the "extras" that really were requested verbally, the reputation of an entire trade stained by those who run. This chapter deals with the most expensive contract a family signs while signing almost nothing. The building's levy and communal works are in chapter 37; the professional's collection, in chapter 34; industrial construction, in chapter 94; home insurance, in chapter 28. The key question: how do you argue about a renovation when all there is are WhatsApps, an unsigned PDF and memories? The answer: giving a certain date to every step — the quote, the initial condition, the phases, the payments, the acceptance — with **t2p**: the fingerprint-time-anchor chain described in chapter 13 [3]. It protects the family from the scammer — and the honest professional from the client who does not want to pay.

The contract that never gets written

Far more frequent than the criminal scam is the ordinary conflict, and everyone has heard a version of it: the verbal quote that grows — "things came up"; the promised quality no one can check; the deadlines that stretch without a date; the payments that run ahead of the work done; the "extras" never agreed in writing; the job left half-done with the relationship broken. The framework exists — the Civil Code regulates the works contract, and consumer law requires a prior quote and an invoice when the client is a consumer [2] — but the typical evidence of a home renovation is an unsigned PDF, a WhatsApp thread and transfers with no reference line.

The renovation has every feature of the perfect conflict: a high amount, a long duration, specifications that evolve, two parties with no documentary habits — and an object, the house, that changes every day, erasing the previous state. When the argument arrives, the client cannot prove what was agreed and the professional cannot prove what he did. And worst of all: the dwelling's initial condition — the wall that "was already cracked," the parquet that "was already scratched" — physically disappears beneath the works themselves, forever.

The rights exist; the facts do not. Every message can be deleted, every transfer goes without a reference, every verbal agreement lives in two diverging memories. What is missing is not more regulation: it is that what is agreed and what is done leave a dated trail in the hands of both.

The proof of existence applied

Five minutes on the first day, ten seconds at each milestone [3].

1. The quote and its versions

The accepted quote — and every modification agreed afterward — scanned and sealed with t2p on the day it is agreed: up to 50 pages in a single certified PDF [3]. The "extras" stop being an argument: either they

have a date and a document, or they do not exist. It protects the client from the quote that grows on its own — and the professional from the client who "does not remember" requesting the change.

2. The initial condition: the photo that protects both

Before the first hammer blow, a photographic walk-through of the dwelling with forensic GPS [3]: the walls, the floors, the windows, whatever was already damaged. It is the evidence the works will physically destroy — and the only defense against the two symmetrical accusations: "you ruined my parquet" and "it was already like that." Five minutes worth the entire conflict.

3. The phases with milestones: paying against facts

Each milestone of the works — the demolition finished, the wiring in, the tiling done — documented with photos and a short video with the date and time overlaid [3], and the corresponding payment linked to the milestone in the same project in the app: quote, phase, evidence, receipt. The advance stops being an act of faith: you pay against documented facts — the pattern that disarms the advance scam [1], because whoever does not plan to come back does not want milestones.

4. The signed acceptance

The finished works — or each phase — are accepted with t2p's strong signature: the video selfie with liveness challenges and biometric verification, from both parties [3]. The dated acceptance closes the perimeter of responsibility: defects appearing later are argued against the received, documented state, not against memories. And the entire file — quote, initial condition, phases, payments, acceptance — is exported as an .asice package verifiable offline [3]: for the consumer claim, the civil route or the criminal complaint [1].

The limits, as always, up front: the seal attests documents, conditions and moments — it does not judge the technical quality of the works: that is for the experts and, if necessary, the judge [2]. It does not replace choosing well — references, identifiable professionals, distrust of the impossible quote [1]. And proportionality rules: the plumber who comes for an hour needs no file; the €30,000 renovation does.

The promise and the peril

The promise: families paying against facts rather than promises [1][3]; honest professionals armored against the bad client — and a trade recovering the reputation the scammers steal from it; disputes over extras and deadlines resolved by looking at dated documents [2]; and the advance scam growing ever harder: whoever does not plan to return does not accept milestones.

The perils, which are not few.

The first is **the client-as-dossier**: the one who documents obsessively to build excuses for not paying. The symmetry disarms him: if every phase has evidence and a signed acceptance, the final non-payment is left without a story — the same trail that protects the client binds him.

The second is **friction in the small trade**. The self-employed worker doing three jobs a day cannot manage dossiers. The gesture must be the ten-second one — the milestone photo, the scanned quote [3] — and, on reflection, it is the same gesture he already performs on WhatsApp: only with a certain date and without depending on a chat.

The third is **the scammer who also documents**: the fake professional's theater with pretty photos. The trail does not replace prudence: the scam's pattern comes before the works — the impossible price and the hurry for the advance [1] — and no seal validates a quote; it validates that it existed.

And a fourth: believing that documenting well is building well. No fingerprint makes a job good or a trade honest. Quality is made by hands and materials; the proof only guarantees that, if it goes wrong, the argument will be about facts. And the best rule is still the old one: never pay far ahead of what has been done.

It's time to act

Let trust be earned with milestones, not promises.

For families: the five-minute ritual — sealed quote, photographed initial condition, payments against milestones [3]; and faced with the impossible price and the rush for the advance, the answer is no [1].

For renovators and the trade guilds: the dated file is your best salesman — the quote, the phases and the signed acceptance [3] armor you against the bad client and set you apart from the scammer who stains your trade [1].

For property managers: demand the verifiable file in the works you supervise — the levy of chapter 37 and the private renovation share the same logic.

For the consumer authorities and the Guardia Civil: the victim's .asice file [1][3] is the difference between a shelvable complaint and an investigation with facts.

For legislators: consumer law already requires a prior quote and an invoice [2]; the natural step is to recognize the quote and the acceptance with a certain date as qualified proof of what was agreed.

May the next half-finished kitchen at least have a complete file — and may the professional who finishes the job collect his pay without arguing over memories. For the homes being renovated and the trades that raise them, **time2act**.



Figure 38.1 — The renovation with proof of existence (sealed with t2p — chapter 13): the quote and its versions sealed on the day of agreement, the initial condition photographed with forensic GPS (the evidence the works erase), the phases with milestones — paying against documented facts — and the acceptance signed with a video selfie by both parties. Symmetry: it protects the family from the advance scam and the honest professional from the client who will not pay. Context: Guardia Civil investigations (Vigo: ≥6 victims of up to €10,000; losses >€77,000 in one case). Limits: quality is for the experts; proportionality; never pay ahead of the facts. Sources: press/Guardia Civil; consumer law.

---

[1] Atlántico — *Investigan varias estafas en reformas de viviendas en Vigo y su comarca* (at least six victims; up to €10,000 each; a builder who starts works and disappears); Liberal de Castilla — *Un vecino de Guadalajara pierde sus ahorros ante un falso albañil de internet* (Guardia Civil; scammer located in Madrid); Telemadrid — *Destapamos al presunto estafador de las reformas*; Allende Abogados — *Estafas en reformas* (pattern: attractive digital ad, cheap quote, advance "for materials," disappearance; losses exceeding €77,000 in some cases). https://www.atlantico.net/articulo/vigo/investigan-estafa-reformas-menos-media-docena-victimas/20221017222637946245.html · https://www.liberaldecastilla.com/un-vecino-de-guadalajara-pierde-sus-ahorros-ante-un-falso-albanil-de-internet-la-guardia-civil-acorrala-al-estafador-en-madrid/ · https://www.telemadrid.es/programas/la-redaccion/Destapamos-presunto-estafador-reformas-2-2339786043--20210510052540.html · https://www.allendeabogados.com/estafas-reformas/ ↩

[2] Framework of the works contract (Civil Code) and consumer law applicable to renovations for consumers (prior quote and invoice). See chapter 34 for the professional's collection and chapter 37 for communal works. ↩

[3] t2p — product manual (document scanner up to 50 pages in one certified PDF; photo with forensic GPS — genuine fix ~50 m / ≤30 s; video with overlaid date/time; projects to link quote, phases, evidence and payments; signature with video selfie, liveness challenges and biometrics; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the levy, in chapter 37; the invoice, in chapter 34. ↩

Real-World Assets: Tokenizing With Proof, Not Blind

"Every man has a property in his own person... The labour of his body, and the work of his hands, we may say, are properly his."

— John Locke, Second Treatise of Government, V, §27 (1689)

On March 20, 2024, the world's largest asset manager made a move that shook Wall Street. BlackRock launched its first tokenized fund, BUIDL, on the public Ethereum network: US Treasury bills, cash and repos placed inside a token that can move across a blockchain. Within twelve months, the fund passed **one billion dollars** [4][5]. Its chief executive, **Larry Fink**, had announced it in his annual letter to investors as "the next generation for markets": turning stocks, bonds and real estate into tradable digital tokens capable of "democratizing yield."

But that move carried a detail that holds all this chapter's tension. The token lives on the chain; the real bonds that back it do not. They remain with a custodian, **off-chain**, and the token merely *points* to them. Fink himself admitted it with revealing candor: for this to scale, he said, identity verification must be solved — knowing who is on the other side and that the token is genuinely backed [10]. That is this chapter's question: when we tokenize the real world, does the token point to a **promise** or to a **proof**?

The great bet — and its hole

The economic promise is dizzying, though nobody agrees on the figure. The consultancy **BCG**, with the private-markets exchange ADDX, projected in 2022 in its report *Relevance of On-chain Asset Tokenization* that the market would grow fiftyfold, to some **$16.1 trillion by 2030** — 10% of world GDP [1]. **McKinsey**, as we have seen, foresaw some **$2 trillion** by 2030, with a range of 1 to 4 [2]; the difference — eightfold — is no technical detail: it reflects how much garbage each is willing to admit and the real speed of adoption. Standard Chartered projects some **$30.1 trillion** by 2034 [9].

The current reality is more modest and more revealing. By mid-2025, the tokenized RWA market — excluding stablecoins — stood at some **$24 billion**, up 380% in three years [3]. The largest segment was not Treasury bonds but **private credit** (some $14 billion), followed by Treasuries (about $7,000M) and tokenized gold (about $4,000M). That gold illustrates both problem and solution well: the PAXG token is backed by serial-numbered bars stored in Brink's vaults in London, and since February 2025 the auditor KPMG has issued monthly attestations verifying the backing ounce by ounce [6]. The question is: what if the attestation could be faked or backdated? The whole edifice rests, once again, on a proof.

The oracle problem

Here lies the structural fragility of all tokenization, and it has a technical name: the **oracle problem**. The token is on the chain, but the value is off it — the gold bar, the deed, the bond, the custody report — and someone has to bridge the two worlds. The **Financial Stability Board** (FSB), the body that coordinates global financial regulation, warned about it plainly in its report *The Financial Stability Implications of Tokenisation* (October 2024): tokenization "often depends on external providers, such as custodians to hold the reference assets and oracles to collect their data," which "introduces new points of vulnerability" and can produce a "mismatch between the token's price and the value of the reference asset" [7]. The rating agency **Moody's** added that many issuers and technology providers have a "limited track record" and that an oracle failure can interrupt payments even when the code works perfectly [8].

This is no abstract fear. In March 2022, the Cashio stablecoin, on the Solana network, collapsed when an attacker exploited a bug that let him mint "backed" tokens... with no real collateral at all [8]. It is the literal image of the principle running through this book: "garbage in, garbage on-chain." A token perfectly valid on the chain that was backed by nothing. And here it is worth recalling the lesson already drawn with carbon credits: when you tokenize on an unverified underlying, the chain fills with "zombie credits." Transposed to real-world assets, the same mistake would repeat at trillion scale.

Tokenizing blind or tokenizing with proof

The conclusion is sharp, and it is the chapter's thesis. Blockchain does not make what it tokenizes true; it makes it *tradable and immutable*. If what lies beneath is a promise without proof, we will have an immutable promise — which is exactly the worst outcome: a lie sold with the appearance of cryptographic certainty. The difference between a solid $2 trillion market and an inflated $16 trillion one is drawn not by the token's technology but by **the quality of the proof of what lies beneath**. Tokenizing blind is repeating the zombies' mistake; tokenizing with proof is the condition for this financial revolution not to end in a crisis of confidence.

The proof of existence beneath the token

This is where **t2p** supplies the missing piece. It does not solve custody or legality — that belongs to the custodian, the auditor and the registry — but it **seals the documentary evidence a token should reference**, verifiably and offline. The mechanism is the usual one, applied to the backing. The document constituting the token's backing — the deed, the appraisal report, the custody or audit certificate, the policy, the photograph of the bar with its serial number — receives a **SHA-512 fingerprint** computed on the device, is anchored to the **public blockchain** and is accompanied by an **RFC 3161 timestamp** conforming to eIDAS. For physical assets, the capture carries **forensic GPS** and a sealed date: "this photo of the asset was taken here, on this day." And the **proof of life** ties the evidence to a real person — the appraiser, the custodian, the signer — resolving precisely the identity verification Fink flagged as the bottleneck.

Everything is packaged into an **.asice container verifiable offline**, independent of t2p and of the token's chain. And here is the turn: the RWA token can **point to the fingerprint of an .asice** instead of pointing to a promise. The buyer stops trusting that there is a bond behind it; they can **check for themselves** that the evidence existed intact at that moment, place and person. This closes the door on substituting or post-editing the backing document, on reusing the same report for several tokens, and on opacity about *when* the proof existed. Turning "trust that there is an asset" into "here is the seal of the evidence — verify it" is, in itself, a transformation of the risk.

One must be, as always, scrupulously honest about the limit, because forgetting it here is especially dangerous. The proof of existence demonstrates that **a record existed, intact, at a moment, place and person**; it does **not** demonstrate that the bar is still in the vault, that it has not been pledged twice, or that the token confers a legal right over it. Custody, independent inventory audits and the legal enforceability of the right remain indispensable. t2p reinforces the **evidence layer**; it does not replace the custodian or the auditor; it is complementary, for instance, to proof-of-reserve systems that verify backing by connecting to the custodians. It is, in one sentence, the difference between tokenizing blind and tokenizing with proof.

The promise and the peril

The promise is immense: tokenization resting on verifiable evidence could open today's illiquid assets — real estate, private funds, infrastructure — to small investors, make markets faster and more transparent, and do so without sacrificing trust. In a market that could be worth trillions, having every token point to a proof rather than a promise is the difference between a revolution and a bubble.

The peril calls for the usual nuance, in the duality of "the promise and the peril." The first, already stated: confusing the seal of the evidence with the guarantee of the legal right or of custody would be a fraud worse than the one being avoided. The second, the false sense of security a brilliant technical layer can lend a bad asset. And the third, concentration: if verification of the evidence rests in very few hands, we recreate the single intermediary blockchain promised to dissolve. The proof of existence is a necessary layer within an architecture that always includes the custodian, the auditor and the regulator.

It's time to act

For **issuers and tokenizers**: anchor the backing's evidence — appraisals, custody certificates, serialized photos — so that every token points to a verifiable proof, not a promise. For **custodians and auditors**: seal your attestations at the origin, so they can be neither backdated nor reused. For **investors**: demand the ability to check the underlying's evidence before buying, and distrust the token that offers only a

promise. For **regulators** (the FSB, national supervisors): recognize the value of the evidence layer and close the oracle gap before the market reaches the trillions being announced.

BlackRock put Treasury bonds inside a token, but the bonds remain off the chain. All real-world tokenization hangs on that invisible thread between the token and the proof of what backs it. Making that thread verifiable by anyone — tokenizing with proof, not blind — is, here too, time2act.

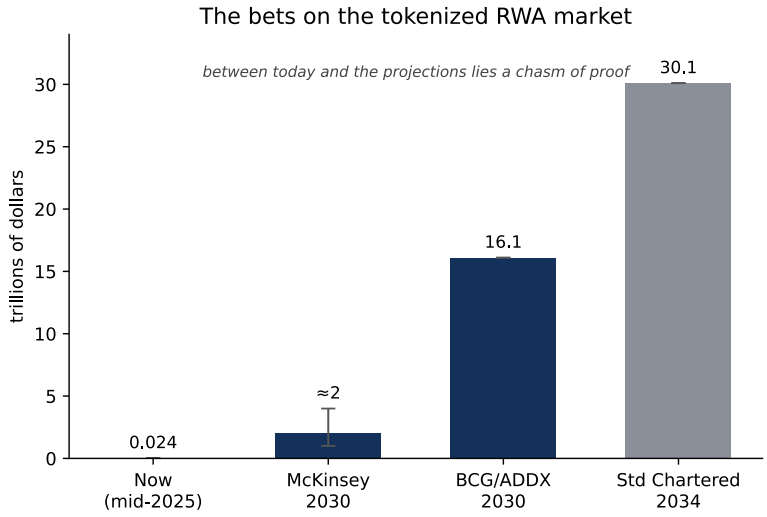


Figure 39.1 — The two reference projections for the tokenized RWA market in 2030 differ about eightfold: McKinsey's prudent scenario (~$2 trillion, range 1-4) and the optimistic BCG/ADDX one (~$16.1 trillion); the real starting point was about $24 billion in mid-2025. The distance depends, at bottom, on how much "garbage" is admitted: on the gap between the token and the proof of its backing — the oracle problem — which the proof of existence helps to close. Sources: McKinsey (2024); BCG/ADDX (2022); CoinDesk/RWA.xyz (2025).

---

[1] BCG and ADDX — *Relevance of On-chain Asset Tokenization in 'Crypto Winter'* (2022; ~$16.1 trillion by 2030, ~10% of world GDP). https://www.bcg.com/publications/2022/relevance-of-on-chain-asset-tokenization ↩

[2] McKinsey — *From ripples to waves: the transformational power of tokenizing assets* (2024; base case ~$2 trillion, range 1-4, excluding stablecoins). https://www.mckinsey.com/industries/financial-services/our-insights/from-ripples-to-waves-the-transformational-power-of-tokenizing-assets ↩

[3] CoinDesk (RedStone / Gauntlet / RWA.xyz report) — *Real-world asset tokenization market has grown almost fivefold in 3 years* (~$24 billion; private credit ~$14 billion). https://www.coindesk.com/business/2025/06/26/real-world-asset-tokenization-market-has-grown-almost-fivefold-in-3-years ↩

[4] BlackRock / BusinessWire — *BlackRock Launches Its First Tokenized Fund, BUIDL, on the Ethereum Network* (20-Mar-2024). https://www.businesswire.com/news/home/20240320771318/en/BlackRock-Launches-Its-First-Tokenized-Fund-BUIDL-on-the-Ethereum-Network ↩

[5] Securitize / PR Newswire — *BlackRock's BUIDL surpasses $1B in AUM* (March 2025). https://www.prnewswire.com/news-releases/blackrock-usd-institutional-digital-liquidity-fund-buidl-tokenized-by-securitize-surpasses-1b-in-aum-302401480.html ↩

[6] Paxos — *PAXG transparency* (serialized bars at Brink's London; monthly KPMG attestations). https://www.paxos.com/paxg-transparency ↩

[7] Financial Stability Board — *The Financial Stability Implications of Tokenisation* (October 2024; custodians and oracles as new points of vulnerability). https://www.fsb.org/2024/10/the-financial-stability-implications-of-tokenisation/ ↩

[8] Moody's (via CoinDesk) — *Tokenized fund adoption grows but brings technology risks* (2024; oracle failures, limited track record). https://www.coindesk.com/markets/2024/01/15/tokenized-fund-adoption-grows-but-brings-technology-risks-moodys ↩

[9] Standard Chartered and Synpulse — *Trade finance to play substantial role in $30.1 trillion tokenised real-world assets market by 2034.* https://www.sc.com/en/press-release/trade-finance-to-play-substantial-role-in-usd-30-1-trillion-tokenised-real-world-assets-market-by-2034/ ↩

[10] InvestmentNews — *Larry Fink takes aim at investment inequality* (BlackRock's 2025 annual letter; identity verification as the condition). https://www.investmentnews.com/alternatives/larry-fink-takes-aim-at-investment-inequality/259910 ↩

Creative Authorship: Proving Who Did It First in the Age of AI

"For masterpieces are not single and solitary births; they are the outcome of many years of thinking in common [...] so that the experience of the mass is behind the single voice."

— Virginia Woolf, A Room of One's Own (1929)

On April 4, 2023, an anonymous user calling himself *ghostwriter977* uploaded a song to the streaming platforms. It was titled "Heart on My Sleeve" and it sounded like an unreleased collaboration between Drake and The Weeknd, two of the biggest artists on the planet. It was not: their voices had been cloned with artificial intelligence. The song spread like wildfire — on TikTok alone it racked up some **fifteen million views** — until, the following Monday, platform by platform, it disappeared. On YouTube the notice remained: "This video is no longer available due to a copyright claim by Universal Music Group" [7]. The label got it taken down by arguing that training and generating AI with its artists' voices violated its agreements and the law.

The scene contains all this chapter's unease. Two of the world's most powerful musicians were left, for a few days, at the mercy of an anonymous figure, and the song did not even contain their recordings: only their synthetic voices. For a creator without a multinational's legal army — an illustrator, a photographer, a screenwriter, an independent musician — the question is more anguishing still: how do I prove that **a work is mine and that it existed before** a machine swallowed it or someone copied it? That is what this chapter is about: making authorship provable again in the age of the infinite copy.

The machinery that swallows authorship

The song's case is the tip of an iceberg of litigation that is redefining creation. The **New York Times** sued OpenAI and Microsoft in December 2023 for copying "millions" of its articles without permission or payment to train their models; as evidence, it attached a 127-page exhibit lining up, column against column, paragraphs from the paper and near-identical outputs from GPT-4 — the visual demonstration that the model "memorized" the journalism it had supposedly only "learned" [1]. The case grinds forward: in January 2026, a judge ordered OpenAI to hand the plaintiffs a sample of twenty million ChatGPT conversation logs [2]. **Getty Images** took Stability AI to court in London for processing millions of its images without a license [3][4]. And a group of artists sued Stability, Midjourney and DeviantArt for feeding the generators with their work.

One of those artists is **Karla Ortiz**, an illustrator who has shaped Marvel characters. In July 2023 she told the US Senate Judiciary Committee a sentence that sums up a whole profession's grievance: regarding practically the entirety of her work ingested by the models, "I was never asked, never credited, never compensated a single cent" [6]. This is not a rejection of the technology. As Getty itself stresses — and it does license its catalog to AI developers while respecting rights — the conflict is not AI as such, but **consent and licensing** [3]. And in Hollywood, the SAG-AFTRA actors' strike, the longest in history against the studios, ended with an agreement requiring informed consent and fair compensation before creating a performer's "digital replica" [8].

The creator and the proof of priority

Beneath all these cases lies one and the same problem, and it is one of **temporal priority**: proving that a song, an illustration, a script or a photo existed, signed by you, *before* a model ingested it or someone copied it. And here the law has given a decisive clue. The US Copyright Office established in January 2025 that **human authorship is the foundation of copyright**: a work generated entirely by AI is not registrable, and the mere choice of prompts, however detailed, is not enough to create a protectable work [10]. The consequence is profound: in a world flooded with synthetic content, value shifts toward **who created what, and when**. The proof of human authorship and of its date becomes the central asset.

Why the label alone gets erased

The industry has already tried to respond with a good idea: the **Content Credentials** of the **C2PA** standard, a coalition under the Linux Foundation that includes Adobe, the BBC, Google, Microsoft, OpenAI and Sony. They work like a "nutrition label" for digital content: a cryptographically signed manifest embedded in the file that records who created it, when, with what tools and with what editing history [9]. It is a valuable, necessary tool. But it has a confirmed Achilles' heel, and it is devastating: **most social networks recompress images on upload and erase the manifest**. WhatsApp, iMessage and Facebook rewrite the file and strip the credentials; a screenshot, a format change or a rescale is enough to break them. C2PA itself acknowledges this and has introduced "durable credentials" with an invisible watermark, but admits the system "depends on compliance across the whole ecosystem" — hardly realistic in an open environment. The label, then, travels inside the file and is lost with it. An independent anchor is missing.

The proof of existence: sealing the work at birth

This is where **t2p** supplies exactly that anchor. Instead of writing the proof *inside* the file — from which it can be erased — it fixes it *outside*, on a public chain. At the moment of creation, the work's **SHA-512 fingerprint** is computed — the demo's `.wav`, the layered `.psd`, the photo's `.raw`, the script — **on the device itself, without the file leaving it**: a crucial detail for the creator who has not yet published and does not want to upload an unreleased work to any server. That fingerprint is anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS, creating a **certain date of authorship** verifiable by anyone, which survives even if the creator changes platforms or disappears. The **forensic GPS** and the **proof of life** tie the moment of creation to a place and to a living person — useful against "who did it first" disputes and against synthetic replicas, because they bind the work to a real human author. And everything is packaged into an **.asice container verifiable offline** that a court or a buyer can check on their own.

The key is that t2p **does not compete with C2PA; it completes it**. Content Credentials provide the readable traceability inside the file; the proof of existence provides the **immutable anchor** that survives even if the manifest is erased or someone takes a screenshot: the hash on the chain still proves that that exact file existed on such a date, place and person. The natural pairing is this: credentials for traceability, proof of existence for the anchor. And the **NFT marketplace** adds a history of provenance and licensing that cannot be edited after the fact, useful for licensing legitimately to those who do want to train with consent — the Getty model: license, not plunder.

The limits must be stated without hedging, faithful to the duality of "the promise and the peril." "Garbage in, garbage on-chain": the proof of existence demonstrates **when, what, where and who** sealed — not that the work is **original or non-infringing**. If you seal a plagiarism, you will have a plagiarism with a certain date. Nor does it settle the underlying legal question — whether training on protected work is fair use or not — which courts and regulators will decide. And the human authorship the law requires remains a separate requirement: the seal proves precedence; it does not make registrable what the law does not protect. Finally, anchoring before publishing only helps if it happens before the leak; it does not rewind what has already been plundered.

The promise and the peril

The promise is immense for millions of creators. A layer of authorship proof at the origin would give them, for the first time, a cheap and universal way to demonstrate their work's priority, to defend themselves in court and to license it on their terms instead of watching it plundered. In a creative economy threatened by the infinite copy, being able to say "this is mine, and I can prove it" is recovering the power to negotiate.

The peril calls for nuance. First, not confusing precedence with originality: the seal does not launder a plagiarism or by itself settle legal authorship. Second, the burden on less technological creators, which the design must make trivially easy so as to exclude no one. And third, privacy: the proof of life and the geolocation of the creative act are sensitive data that must be protected. The proof of existence is a tool of creator empowerment, not a shortcut that replaces the courts or the law.

It's time to act

For **creators** (musicians, illustrators, photographers, screenwriters): seal the work at the moment of creation, before sharing it, to hold a certain date of authorship. For **platforms**: stop destroying provenance credentials on upload and recognize the independent anchor. For **AI developers**: license work with consent — the Getty model — instead of plundering it, using the provenance trail to identify and compensate authors. For **regulators and courts**: close the legal framework of consent and fair use, and recognize the evidentiary value of the authorship seal.

Two of the biggest musicians in the world were left defenseless before a synthetic voice; the individual creator is more defenseless still. The difference between a work anyone can claim or copy and a work born with an unalterable proof of authorship is, increasingly, the difference between creating and being plundered. As we have seen, proving who did it first — and that they were human — is, here too, time2act.

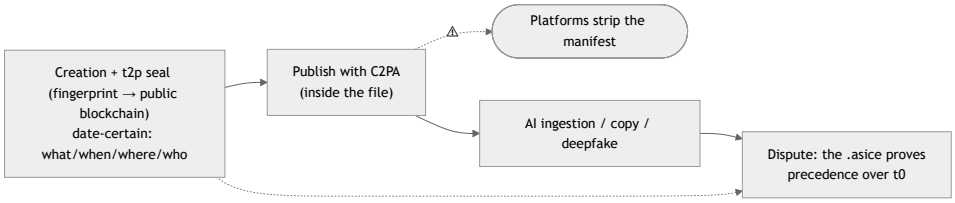


Figure 40.1 — At the moment of creation (t0), the work is sealed with a SHA-512 fingerprint on the device, an anchor to the public blockchain, an eIDAS timestamp, GPS and proof of life: a certain date of what, when, where and who. On publication (t1), the C2PA credentials can be erased on upload to the networks; when a model ingests it or someone copies it (t2), the .asice container still proves the file existed intact at t0, before t2. The proof of existence proves precedence, not originality. Sources: C2PA; U.S. Copyright Office; the NYT, Getty and Andersen cases.

---

[1] NPR — *The New York Times sues OpenAI and Microsoft for copyright infringement* (27-Dec-2023; 127-page exhibit with GPT-4 outputs). https://www.npr.org/2023/12/27/1221821750/new-york-times-sues-chatgpt-openai-microsoft-for-copyright-infringement ↩

[2] Bloomberg Law — *OpenAI must turn over 20 million ChatGPT logs, judge affirms* (January 2026). https://news.bloomberglaw.com/ip-law/openai-must-turn-over-20-million-chatgpt-logs-judge-affirms ↩

[3] Getty Images — *Getty Images statement* (action against Stability AI; consent and licensing, January 2023). https://newsroom.gettyimages.com/en/getty-images/getty-images-statement ↩

[4] Judiciary of England and Wales — *Getty Images v Stability AI [2025] EWHC 2863 (Ch)* (Nov 2025; dismisses most of the copyright claims). https://www.judiciary.uk/judgments/getty-images-v-stability-ai/ ↩

[5] Copyright Alliance — *Andersen v. Stability AI* (Sarah Andersen, Kelly McKernan, Karla Ortiz; discovery since August 2024). https://copyrightalliance.org/andersen-v-stability-ai-copyright-case/

[6] Karla Ortiz — *Testimony before the U.S. Senate Judiciary Committee* (12-Jul-2023; "never asked, never credited, never compensated"). https://www.judiciary.senate.gov/imo/media/doc/2023-07-12_pm_-testimony-_ortiz.pdf ↩

[7] Variety — *AI-generated fake Drake/The Weeknd song 'Heart on My Sleeve'* (April 2023; pulled by Universal Music Group). https://variety.com/2023/music/news/fake-ai-generated-drake-weeknd-collaboration-heart-on-my-sleeve-1235585451/ ↩

[8] Wikipedia — *2023 SAG-AFTRA strike* (14-Jul to 9-Nov-2023; consent for the "digital replica"). https://en.wikipedia.org/wiki/2023_SAG-AFTRA_strike ↩

[9] C2PA / Content Authenticity Initiative — *Content Credentials* (signed manifest; platforms erase it on recompression). https://c2pa.org/ ↩

[10] U.S. Copyright Office — *Copyright and Artificial Intelligence, Part 2: Copyrightability* (29-Jan-2025; human authorship as the foundation). https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf ↩

Patents and Prior Art: Defensive Publication With a Certain Date

"One can resist the invasion of armies; one cannot resist the invasion of ideas."

— Victor Hugo, The History of a Crime (1877)

There is a scene that repeats itself in patent attorneys' offices all over the world. An independent inventor arrives with a sealed envelope, yellowed by the years. Inside is the description of his invention, which he mailed to himself by registered post long ago, convinced that the postmark was a kind of **"poor man's patent"** that would protect his idea. Now a competitor has patented something similar, and the inventor wants to assert his envelope. The attorney has to explain, tactfully, that this gesture repeated for decades **never protected anything**: mailing yourself an envelope establishes no patent right and is not evidence recognized by the courts [5]. And there is worse news still: since 2013, in the United States, not even the date matters against whoever filed the application first.

This scene contains the knot of this chapter. Industrial property is an expensive, relentless race, and most inventors — the self-employed, the small companies, the researchers — cannot afford to run in it. But there is a powerful, little-known move that changes the rules for them: instead of patenting, **publishing with a certain date** to prevent anyone else from patenting it. And here the proof of existence offers an almost tailor-made tool.

The race of whoever files first

To understand the move, you must know the terrain. On March 16, 2013, the United States abandoned the old "first to invent" system and adopted **"first to file"** (*first-inventor-to-file*), with the entry into force of the America Invents Act. As the US patent office itself, the **USPTO**, sums it up, from that date on "the date of invention is irrelevant": what counts is who files the application first [1]. The system rewards speed and capital, not the paternity of the idea.

According to industry estimates, drafting a utility patent application in the United States costs between $9,000 and $17,000, and one must budget **between $15,000 and $25,000** before the patent is granted, not counting the maintenance fees that stretch across its whole life [3]. For a large corporation that is small change; for an individual inventor, an often insurmountable barrier. There is a cheaper official route to fix the date — the provisional application, which for a few hundred dollars grants "patent pending" status for twelve months — but it is not examined and lapses if the full application is not filed afterward [4].

One key concept opens a door. The **World Intellectual Property Organization** (WIPO) defines the **prior art** as all information made available to the public anywhere in the world before the filing date; for an invention to be patentable, it must contribute a new characteristic not found there [2]. And here comes the decisive turn: WIPO reminds us that **publicly disclosing an invention before filing can destroy its novelty** and render it unpatentable — unless the law provides a "grace period," present in the United States or Canada but not in Europe, and in China only in very restricted form [2]. This double-edged sword is precisely the defensive weapon.

Publishing to defend yourself

If disclosure destroys novelty, you can disclose on purpose to keep someone else from patenting. This is the **defensive publication**: deliberately disseminating a description of the invention so that it enters the public domain and becomes prior art, defeating any later patent application by a third party on the same idea. It grants no exclusive rights — once published, the invention can no longer be patented — but it avoids the costs of a patent and, above all, prevents a competitor from taking it from you and charging you for its use.

This is no marginal idea; it is an industry. **Technical Disclosure Commons**, a **free** platform operated by Google, lets anyone publish a technical disclosure that the system **dates automatically** and that **Google Patents indexes**, so an examiner will find it when evaluating a rival application [6]. **IP.com** offers a prior art database where every publication is closed with a **cryptographic timestamp** — from the Surety AbsoluteProof service — protecting its integrity and its date. And **Research Disclosure**, the

sector's oldest service, has published defensive disclosures since 1960. All this material is searchable in gigantic databases: the European Patent Office's Espacenet passed **160 million documents** in April 2025, and WIPO's PATENTSCOPE hovers around 128 million [9][10]. As the patent attorneys Richard Coller and Alexander Covington explain in a reference piece on the strategy, defensive publication never replaces the value of a patent for an important innovation, but it is a highly cost-effective way to block competitors when hard budget decisions must be made [8].

Note one detail, because it is the key to everything: IP.com already seals its publications with a cryptographic timestamp and an archive anyone can validate. It is, in essence, a **digital notary** that freezes content and date. And this is where t2p comes in, generalizing that service to anyone.

The proof of existence for the invention

The mechanics of **t2p**, applied to the invention, are simple and potent. The inventor takes the disclosure document — the technical memorandum, the drawings, the code, the laboratory notebook where conception is recorded — and computes its **SHA-512 fingerprint** on the device, without the document leaving it. That fingerprint is anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS, and everything is packaged into an **.asice container verifiable offline**. The **forensic GPS** and the **proof of life** can tie the moment to a place and a real person.

The result proves, unalterably and verifiably worldwide, that **that exact content existed intact on that date**, optionally in that place and by that person. It is, in practice, two things at once: an **instant, cheap defensive publication with a certain date** — if, in addition, the content is made public on that date — and a **proof of the date of conception** for disputes between co-inventors, between company and employee or between partners. Against the fifteen-to-twenty-five-thousand-dollar budget of a patent, the cost of a seal is marginal. And against backdating it is impenetrable: no one can later fabricate a document "from that date," because the anchored fingerprint and the seal are immutable.

The limits must be stated in all their harshness, faithful to the duality of "the promise and the peril," because getting confused here is expensive. The proof of existence **does not prove novelty, patentability, utility or that the invention works**: the examiner decides that against the whole world's prior art. Nor does it turn a private document into prior art by itself: to block a third party, the disclosure must be **accessible to the public** — an anchored hash without publishing the content proves *possession on a date*, not *availability to the public*. And in Europe, where there is no grace period, publishing **also destroys your own novelty**: the certain date does not replace the strategic decision between patenting and publishing; it only makes it more solid once taken. "Garbage in, garbage on-chain": the seal proves date and content, not that the idea is genuinely yours or original.

The promise and the peril

The promise is deeply democratizing. A cheap certain-date tool puts within the individual inventor's reach a move — defensive publication — that until now was dominated by large corporations with patent departments. It makes it possible to protect an idea's priority, block a giant from patenting it and settle authorship disputes, all for a fraction of the system's cost. In an innovation economy where speed and capital decide everything, it means returning a little power to those who create.

The peril calls for nuance. First, as stated: confusing the seal with a patent or with proof of novelty would be a costly error; the proof of existence is precedence, not patentability. Second, the temptation to seal without publishing and believe yourself protected against third parties, when only public disclosure creates opposable prior art. And third, the irreversible decision: publishing defensively closes forever the door to patenting that invention. The proof of existence is a tool in the service of a strategy, not a substitute for expert advice on when to patent and when to publish.

It's time to act

For **individual inventors and small companies**: seal the conception with a certain date and, when the move is defensive, publish the disclosure to block others' patents at a fraction of the cost. For **R&D departments**: integrate sealing into the workflow, to hold proof of priority for every idea before deciding what gets patented and what gets published. For the **defensive publication platforms**: offer the

independent anchor as a complement to your seal. For the **patent offices**: recognize and index these dated disclosures as prior art in their own right.

The inventor with the yellowed envelope entrusted his idea to a gesture that protected nothing. Today, anyone can leave an unalterable record that an idea was theirs on a certain date and turn it into a wall against whoever wants to patent it. Proving who conceived it first — and putting that within everyone's reach — is, here too, **time2act**.



Figure 41.1 — The inventor seals the disclosure document with a SHA-512 fingerprint, an anchor to the public blockchain and an eIDAS timestamp (certain date, integrity and authorship), at a marginal cost against the $15,000-25,000 of a patent. If they also publish it (Technical Disclosure Commons, IP.com, the web), it becomes prior art indexed by Espacenet or Google Patents; when a rival's application arrives, the examiner finds the dated disclosure and rejects it for lack of novelty. The seal proves date and content, not patentability — which the examiner decides. Sources: USPTO; WIPO; Technical Disclosure Commons.

---

[1] USPTO — *America Invents Act (AIA): Frequently Asked Questions* (first to file since 16-Mar-2013; "the date of invention is irrelevant"). https://www.uspto.gov/patents/laws/america-invents-act-aia/america-invents-act-aia-frequently-asked ↩

[2] WIPO (World Intellectual Property Organization) — *Patents: Frequently Asked Questions* (prior art; prior disclosure destroys novelty; the grace period is not universal). https://www.wipo.int/en/web/patents/faq_patents ↩

[3] BitLaw — *What does a patent application cost?* (drafting $9,000-17,000; budget $15,000-25,000 before grant). https://www.bitlaw.com/guidance/patent/what-does-a-patent-application-cost.html ↩

[4] USPTO — *Fee Schedule* (provisional application: fixes the date, 12 months; official fees). https://www.uspto.gov/sites/default/files/documents/USPTO-fee-schedule_current.pdf ↩

[5] McLane Middleton — *Know the Law: Is the 'Poor Man's Patent' Real?* (no legal value; obsolete since the 2013 first-to-file). https://www.mclane.com/insights/know-the-law-is-the-poor-mans-patent-real/ ↩

[6] Technical Disclosure Commons — *FAQ* (free; automatic publication date; indexed by Google Patents). https://www.tdcommons.org/faq.html ↩

[7] IP.com — *Publishing Disclosures to the Prior Art Database* (timestamping via Surety AbsoluteProof). https://kb.ip.com/pad/knowledge-base/publishing-disclosures-to-the-prior-art-database-pad/

[8] Richard Coller III and Alexander Covington (IPWatchdog) — *Defensive Publications: A Cost-Effective Tool to Supplement Your Patent Strategy* (2020). https://ipwatchdog.com/2020/05/25/defensive-publications-cost-effective-tool-supplement-patent-strategy/ ↩

[9] European Patent Office (EPO) — *Espacenet reaches 160 million patent documents* (April 2025). https://agip-news.com/news.aspx?id=76074&lang=en ↩

[10] WIPO — *PATENTSCOPE* (~128 million patent documents; 5.4 million PCT applications). https://patentscope2.wipo.int/ ↩

Trade Secrets: Proving Without Revealing

"No man can be secret, except he give himself a little scope of dissimulation; which is, as it were, but the skirts or train of secrecy."

— Francis Bacon, Essays, "Of Simulation and Dissimulation" (1625)

One of the star engineers of Google's self-driving car project — the one that would become Waymo — downloads thousands of confidential files onto his laptop before leaving. He founds an autonomous truck company, Otto. Uber buys it within months. And Waymo sues Uber: years and fortunes of research are now driving in the next lane. In 2018, Uber closes the matter with a settlement the press valued at some $245 million in shares [1]. In March 2020, Anthony Levandowski pleads guilty to trade secret theft; in August, Judge William Alsup sentences him to eighteen months in prison with a phrase for the annals: "the biggest trade secret crime I have ever seen." And the novelistic epilogue: on January 20, 2021, the last day of his term, President Trump pardoned him [1].

If this happened to Google — with the best lawyers and systems in the world — what does not happen every week at small scale: the salesman who leaves with the client list, the partner who walks out with the factory procedure, the supplier who "reinvents" the design you showed him under a confidentiality agreement. And the victim then discovers the knot of the trade secret: to defend it before a judge you must prove what you had, since when, and that you protected it — without revealing it a second time along the way.

This chapter deals with the only asset that is defended by hiding it — and which for that reason has the worst evidentiary problem in the law. It is the mirror of chapter 41: the patent protects by publishing; the secret, by concealing. Creative authorship is in chapter 40; the logs an attacker cannot rewrite, in chapter 57; and the whistleblower who reveals wrongdoing — who is not a secrets thief — in chapter 62. The key question: how do you prove you have something you cannot show? The answer: with the fingerprint-sealed deposit of **t2p** — the SHA-512 computed on the device, the RFC 3161/eIDAS timestamp, the anchoring to the public blockchain [3] — where only the fingerprint travels and the content never leaves home: possession with a certain date, an exact perimeter, and no disclosure. And with the symmetry that keeps the chapter honest: the same tool defends the former employee or the unjustly accused competitor, who can prove their prior independent development.

Three requirements and one knot

The framework is Law 1/2019 on trade secrets — the transposition of the European directive — and it is demanding [2]: for information to be a protectable secret it must meet three requirements. It must be secret — not generally known or readily accessible in the circles that deal with this kind of information. It must have business value precisely because it is secret. And it must have been subject to reasonable measures by its holder to keep it secret. Chemical formulas, production processes, databases, strategies, client lists, algorithms, know-how: everything fits — if you meet the three requirements.

The evidentiary knot lives in the third. Without demonstrable reasonable measures, there is no protectable secret: the information can be enormously valuable and leaked all the same, and the judge will tell you that you did not protect it enough. The documentation of the measures thus becomes part of the secret itself — and its absence, the classic trap of the careless company.

The other two strands of the knot: to litigate you must delimit exactly what the secret was — and every extra word you reveal in the proceedings is secret that dies, so the drafting must be surgical; and you must prove temporal priority — that you had it first — as with the patents of chapter 41, but with no public registry to record it in: here the date must be procured by the holder himself.

And the balance the law itself imposes — and this book endorses [2]: the protection does not extend to the experience and skills a worker honestly acquires over a career, nor can it serve to restrict mobility. The frontier between "our know-how" and "their trade" is the battle line of almost all these lawsuits — and without dated facts, each side draws it wherever it likes.

The proof of existence applied

Here the proof of existence plays its perfect game, because it is the only one that proves without showing [3].

1. The sealed deposit

The company certifies with t2p the files that constitute its secret — the formula, the code, the client database, the procedure: the SHA-512 fingerprint is computed on the device and only the fingerprint is anchored; the content never leaves home [3]. The result is threefold: possession with a certain date — you had it on day X; an exact perimeter — byte by byte, what is and is not your secret; and a reasonable measure in itself: a dated custody protocol is exactly what the third requirement asks for [2]. On the day of litigation, disclosure is selective and before the judge — the law's procedural confidentiality measures help [2] — and the certificate with its QR code attests that this is what existed on that day.

2. The chronology of the know-how

Every version of the development — the laboratory notebook, the code iterations, the trials — sealed in the moment [3]. The chain of certain dates draws your own path: when the competitor shows up with "the same" product, your chronology proves anteriority and independent development. As in chapter 41, precedence is the heart of the dispute — but here without publishing anything.

3. The employee's departure

The moment of maximum risk, documented with neutrality: the record of the return of equipment and credentials, scanned and sealed — up to 50 pages in one PDF [3]; the inventory of accesses held, anchored; the NDA reminder, delivered with a certain date. No drama, no suspicion: a ten-minute formality that protects both parties — the company, if its files later turn up where they should not; and the former worker, who has a record of what was returned and when.

4. The NDA with a perimeter — and the innocent's defense

The NDA signed with t2p's strong signature — a video selfie with liveness challenges and biometric verification [3] — and, as an annex, the secret's perimeter defined by fingerprint: the recipient knows exactly which files they may not use, without the content traveling. And the symmetry that keeps this chapter honest: the former employee who starts a company, or the competitor who developed their solution on their own, holds in their own chain of seals the best defense against the unfounded accusation [2][3] — their prior knowledge, dated before any contact.

The limits, stated plainly: the seal does not make secret information that is already public, nor does it decide whether something was a protectable secret or professional experience — the judge will say, with the law in hand [2]. It does not replace real security — access controls, encryption, leak prevention: it attests what you had; it does not stop the theft; for the attack's logs, chapter 57. And it can never serve to seal everything and claim everything: the inflated perimeter turns against you in the proceedings.

The promise and the peril

The promise: trade secret lawsuits settled on certain dates and exact perimeters rather than on testimony and suspicion [1][2][3]; the third requirement — reasonable measures — met by design and provable with a certificate; small companies protecting their know-how without paying the price of a procedural war; and workers changing jobs with their baggage documented and their conscience clear.

The perils, faced squarely.

The first is **the hunt for the former employee**. The same file that protects the secret can industrialize the threat against whoever leaves: intimidating letters with sealed annexes to freeze competition. The counterweight is threefold: the law's express balance — honest experience and mobility are untouchable [2] — the proof's symmetry — the defendant's independent development is dated too [3] — and the commercial courts punishing procedural abuse.

The second is **the inflated perimeter**. Sealing everything and calling everything secret does not prove reasonable measures: it proves disorder. The protectable secret is the one that is delimited, guarded and treated as such — the seal documents the discipline; it does not replace it.

The third is **the secret against the public interest**. No sealed deposit turns into a secret what the law obliges to be disclosed, nor can it be used against the whistleblower who reports wrongdoing — that frontier belongs to chapter 62, and it is sacred.

And the fourth is **safe-box techno-solutionism**. The fingerprint does not prevent the leak: it detects and dates it. Whoever confuses the deposit with security will end up with a secret perfectly documented… and perfectly stolen. Encryption, least-privilege access and internal culture remain the first line.

It's time to act

Let what cannot be shown be provable.

For companies and founders: inventory your secrets, delimit them and deposit their fingerprints with a certain date [2][3]; the day you need it, "reasonable measures" will be a certificate, not a plea.

For legal departments and IP advisors: NDAs with a fingerprint perimeter and anchored departure records as standard practice [3]; and remind the client of the law's balance: the secret is defended, the former employee is not persecuted [2].

For the workers and entrepreneurs who leave: seal your prior baggage and your independent development [3] — your professional freedom is worth a chain of dates.

For the commercial courts: the fingerprint deposit with selective disclosure resolves the paradox of the trade secrets trial [2]; demand it — and punish both the theft and the strategic accusation.

For legislators: the 2019 law set the three requirements [2]; the next step is to expressly recognize the fingerprint deposit with a certain date as a reasonable measure and as qualified proof of possession.

May the next Levandowski [1] find fingerprints where he expected fog — and may the next innocent find dates where he expected suspicion. For the knowledge that is valuable precisely because it cannot be seen, **time2act**.

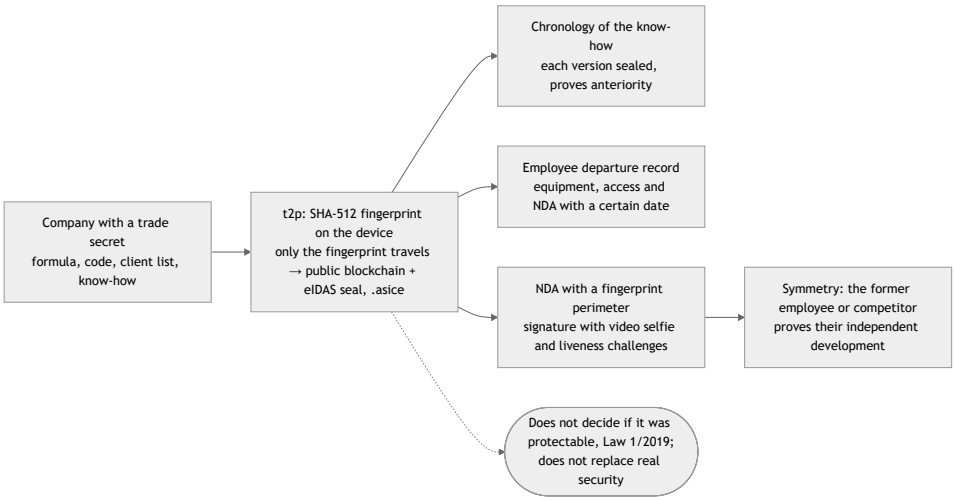


Figure 42.1 — The trade secret with proof of existence (t2p: SHA-512 fingerprint on the device — the content never leaves — RFC 3161/eIDAS, public blockchain, .asice): the sealed deposit (possession + perimeter + the reasonable measure of Law 1/2019's third requirement), the know-how's version-by-version chronology, the employee departure record and the NDA with a fingerprint perimeter and video-selfie signature. Symmetry: the innocent defendant's chain of dates proves independent development. Mirror of chapter 41: the patent protects by publishing; the secret, by concealing — which is why only the proof that does not reveal can defend it. Sources: the Waymo/Uber/Levandowski case; Law 1/2019.

[1] CNBC — *Trump pardons Anthony Levandowski, the engineer who stole self-driving car secrets from Google* (thousands of files downloaded before leaving; Otto bought by Uber; 2018 Waymo-Uber settlement valued by the press at ~$245 million in shares; guilty plea in March 2020; 18 months in prison — Judge William Alsup: "the biggest trade secret crime I have ever seen"; full pardon on 20-1-2021, the last day of the term); TechCrunch — *Former Google engineer Anthony Levandowski among list of last-minute Trump pardons*; Wikipedia — *Anthony Levandowski* (2022 global settlement: $25-30 million). https://www.cnbc.com/2021/01/20/anthony-levandowski-pardoned-after-stealing-trade-secrets-from-google.html · https://techcrunch.com/2021/01/19/former-google-engineer-anthony-levandowski-among-list-of-last-minute-trump-pardons/ · https://en.wikipedia.org/wiki/Anthony_Levandowski ↩

[2] BOE — *Ley 1/2019, de 20 de febrero, de Secretos Empresariales* (transposition of Directive (EU) 2016/943; three requirements: secrecy + business value + reasonable measures; scope: formulas, processes, databases, client lists, algorithms, know-how; balance: it does not cover the experience and skills honestly acquired by the worker and cannot restrict mobility; procedural confidentiality measures); Garrigues — *Aprobada la ley que protege los secretos empresariales*. https://www.boe.es/buscar/act.php?id=BOE-A-2019-2364 · https://www.garrigues.com/en_GB/new/law-will-protect-trade-secrets-has-been-approved ↩

[3] t2p — product manual (certify any file — only the fingerprint travels, the content never leaves the device; document scanner up to 50 pages; signature with video selfie, liveness challenges and biometrics; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package with selective disclosure verifiable offline). See chapters 13 and 15; prior art and the patent, in chapter 41; an attack's logs, in chapter 57; whistleblowers, in chapter 62. ↩

Govtech: Public Procurement and Tamper-Proof Records

"But who will guard the guards themselves?"

— Juvenal, Satires, VI, 347-348 (c. 100 AD)

In June 2020, an unlikely alliance — the World Economic Forum, the Inter-American Development Bank and Colombia's Attorney General's Office — announced an experiment that seemed to step out of a metaphor. They wanted to put on a blockchain the bidding phase of the **School Feeding Program**, which buys meals for the country's most vulnerable children and where, historically, corruption had seeped in. The idea was simple and moving: to make it impossible for anyone to steal, quietly and after the fact, the children's lunch money. The prototype ran on the public Ethereum network and automatically recorded every modification or intervention. But its champions were honest about the limits, and their conclusion holds for this entire chapter: "technology cannot fully resolve what lies at the heart of human behavior" [8].

That experiment touches the nerve of this final front. Public procurement is, as the OECD itself says, the government activity most vulnerable to corruption, because of the volume of money and the discretion it moves. And the trust citizens place in it — that their taxes buy real schools, bridges and meals — rests, once again, on evidence that can be manipulated. This chapter is about making that trust *provable*, and doing so with the humility the Colombian experiment already learned.

The most vulnerable spending

The magnitude of what is at stake is hard to picture. World public procurement exceeds **$13 trillion a year**, close to 13-15% of global GDP; it is, by far, the biggest economic lever in the hands of states [1]. And a chilling share of it evaporates. The **World Bank** calculates that governments waste about **a quarter** of this spending through inefficiency and corruption, and that sound procurement could free up **a million million — one trillion — dollars a year**, what it calls "the hidden trillion" [2]. The World Bank itself tracks the progress of digital government worldwide with its GovTech Maturity Index, covering 198 economies [10]. In the construction sector the hole is more brutal still: the infrastructure transparency initiative **CoST** estimates that **between 10% and 30% of a public work's value** is lost to mismanagement and corruption [3].

A good part of this loss has a mechanism with a name of its own. The OECD, in its guidelines against **bid rigging**, warns that when companies agree to fix a tender, eliminating that collusion could cut contract prices by **20% or more** [4]. The forms are well known: cover bids, bid suppression or rotation, market sharing, often with the complicity of an insider within the contracting body. The global result of this erosion is measured every year by **Transparency International** with its Corruption Perceptions Index: the 2025 edition marked a world average of barely **42 out of 100, an all-time low**, with more than two-thirds of countries failing [5].

Transparency is not the same as proof

Faced with this, states have tried two strategies that should not be confused, because they illustrate the exact difference between *seeing* and *proving*.

The first is **transparency**: opening the data so everyone can look. The best example is **ProZorro**, the open procurement system Ukraine built in 2016 under the motto "everyone sees everything." According to the Open Contracting Partnership, it has saved the treasury some **$6 billion**, and the citizen watchdog community Dozorro has detected irregularities in more than 33,000 tenders [6]. It is a supremely powerful tool — a country at war proved that the anti-corruption weapon was opening the data, not hiding it — but it has a limit: transparency lets you see what is published, but it does not by itself guarantee that the record has not been altered or trick-dated.

The second strategy attacks precisely that limit: **tamper-proof integrity**. The pioneer is Estonia, which after the state cyberattacks of 2007 became, in 2012, the first country in the world to put blockchain into production to armor its records, with the **KSI** technology of the company Guardtime. The design is elegant and is, as we shall see, the same as t2p's: **only the document's fingerprint is sent to the**

chain, never the original data, which does not leave the state's servers [7]. Estonia can thus prove mathematically that no official has retroactively rewritten a medical record, a property entry or a court file. Transparency is letting people see; proof of integrity is being able to prove that what they see has not been touched.

The proof of existence applied to public money

This is where **t2p** delivers Estonia's integrity without any town hall or agency having to build a national blockchain, with three anchoring points.

The first, sealing bids and tender documents against backdating. When a bid is submitted or tender terms published, the document's **SHA-512 fingerprint** is computed on the device, anchored to the **public blockchain** and given a qualified **RFC 3161 timestamp** conforming to eIDAS. This cryptographically fixes that *that exact document existed at that hour*, so no one can antedate a bid, substitute a tender document after the opening or alter an envelope without the fingerprint and the date changing. Moreover, **Article 41 of the eIDAS regulation** grants the qualified timestamp a legal presumption of the date's accuracy and the data's integrity [9], something ProZorro's transparency does not guarantee on its own.

The second, **sealing public records**, Estonian style but as a service: every entry or version of a file — an award, a contract modification, a certificate — is digested into a fingerprint and anchored; any later rewriting becomes mathematically detectable by an external auditor.

And the third, the most distinctive, where public money actually lands: **sealing the delivered work**. An official or a citizen captures the **photograph of the delivered infrastructure** — the bridge, the classroom, the school meal — with **forensic GPS, date and proof of life** of the operator, all packaged into an **.asice container verifiable offline**. An auditor can thus independently confirm that *the money bought something real, in that place, on that date*, striking directly at the loss observed in public works ("as we have seen"). It is the step that closes the circle: from the sealed bid to citizen verification of what was built.

The limit must be stated with the same honesty the Colombian experiment had, faithful to the duality of "the promise and the peril." The proof of existence demonstrates that **a record existed intact at a given moment, place and person**; it does **not** demonstrate the physical truth of what lies behind it or the honesty of the decision. If the sealed bid was already rigged by collusion, if the GPS photo shows a bad work presented as good, or if the official vouching for it is complicit, t2p will faithfully seal a corrupt datum. What it does — and it is no small thing — is **make after-the-fact manipulation and backdating expensive and detectable**, creating an auditable chain of custody. It does not replace the competition investigation, the quality control of the works, or judgment on the merits of the award. "Technology cannot fully resolve human behavior," as they acknowledged in Colombia; but it can close many of the doors through which the money escapes.

The promise and the peril

The promise is measured in that "hidden trillion." Procurement where bids, records and works are sealed and verifiable by any citizen would make backdating, file substitution and the phantom work far harder, and would hand the public a tool to check that their taxes buy what they are said to buy. As we have seen, this sector moves $13 trillion; recovering that trust is, literally, recovering resources for hospitals and schools.

The peril calls for the usual nuance. First, the seal does not launder a corrupt decision or a bad work; confusing the record's integrity with the honesty of the substance would be a dangerous false security. Second, surveillance: an infrastructure that seals every act of the administration must serve to control power, not citizens, and demands a rights-protective design. Third, concentration: whoever operates the verification layer must not become a new single gatekeeper of what is public. The proof of existence is a tool in the service of accountability, not a substitute for honest politics or democratic control.

It's time to act

For **contracting bodies**: seal bids, tender documents and awards at the origin, so the process's temporal integrity is provable and backdating impossible. For **oversight and audit bodies**: demand the seal of

the delivered work — photo, GPS, date — as the basis of scrutiny, to confirm the money acquired something real. For **citizens and journalism**: claim the ability to verify on your own both the open data and its integrity. For **governments and regulators**: combine ProZorro-style data transparency with Estonian-style proof of integrity, and recognize the evidentiary value of the seal that the eIDAS regulation already upholds.

In Colombia they tried to stop anyone from stealing, in silence, the children's food. They discovered that technology does not cure everything — and they were right — but also that making every movement of public money verifiable by anyone is one of the most effective walls against corruption. That a citizen can prove their taxes bought something real is, on this last front as on all the others, **time2act**.

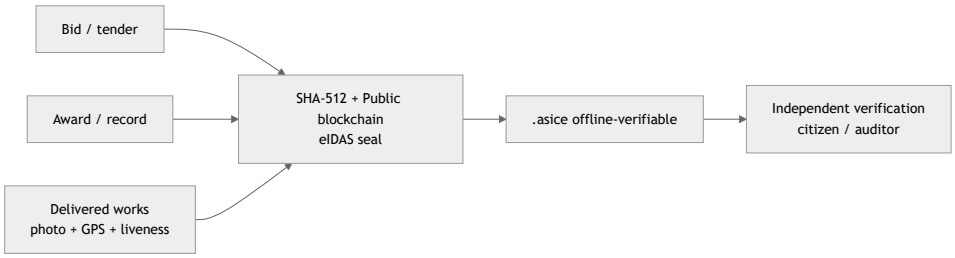


Figure 43.1 — Every step of public money — the bid or tender document, the recorded award and the delivered work (photo with GPS and date) — is digested into a SHA-512 fingerprint, anchored to the public blockchain and given an RFC 3161/eIDAS timestamp, with proof of life at the work site; everything is packaged into an .asice container that a citizen or an auditor can verify independently and offline. Context: world public procurement ≈$13 trillion a year; between 10% and 30% of public works lost to corruption and mismanagement. The proof of existence demonstrates integrity, moment, place and person — not the honesty of the decision or the quality of the work. Sources: World Bank; CoST; OECD; eIDAS Art. 41.

---

[1] Spend Network / World Bank — *The $13 trillion global value of public procurement* (~13-15% of GDP; under 3% in open data). https://www.spendnetwork.com/13-trillion-the-global-value-of-public-procurement/ ↩

[2] World Bank — *The hidden $1 trillion: halting waste in public procurement* (~a quarter wasted; ~$1 trillion/year). https://blogs.worldbank.org/en/voices/hidden-1-trillion-halting-waste-public-procurement ↩

[3] CoST – Infrastructure Transparency Initiative — *This is why construction is so corrupt* (10-30% of public works value; ~$6 trillion/year toward 2030). https://infrastructuretransparency.org/2016/02/05/this-is-why-construction-is-so-corrupt/ ↩

[4] OECD — *Guidelines for Fighting Bid Rigging in Public Procurement (2025 Update)* (collusion can raise contract prices 20% or more). https://www.oecd.org/en/publications/2025/09/oecd-guidelines-for-fighting-bid-rigging-in-public-procurement-2025-update_127880ea.html ↩

[5] Transparency International — *Corruption Perceptions Index 2025* (world average 42/100, all-time low; Denmark 89, eighth year at the top). https://www.transparency.org/en/news/cpi-2025-findings-insights-corruption ↩

[6] Open Contracting Partnership — *The impact of ProZorro in Ukraine* (~$6 billion saved; Dozorro, 33,000+ tenders). https://www.open-contracting.org/impact-stories/impact-ukraine/ ↩

[7] e-Estonia — *KSI Blockchain* (first state with blockchain in production, 2012; Guardtime; only the fingerprint is anchored). https://e-estonia.com/solutions/cyber-security/ksi-blockchain/ ↩

[8] World Economic Forum — *Exploring Blockchain Technology for Government Transparency to Reduce Corruption* (Colombia, School Feeding Program, IDB, 2020). https://www.weforum.org/publications/exploring-blockchain-technology-for-government-transparency-to-reduce-corruption/ ↩

[9] EUR-Lex — *eIDAS Regulation (EU 910/2014), Article 41* (legal presumption of accuracy of the qualified timestamp). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2014.257.01.0073.01.ENG ↩

[10] World Bank — *GovTech Maturity Index* (198 economies; four areas of digital government maturity). https://www.worldbank.org/en/programs/govtech/gtmi ↩

The Citizen Before the Administration: The Clock of Silence

"The concept of the 'official secret' is the specific invention of bureaucracy, and nothing is so fanatically defended by the bureaucracy as this attitude."

— Max Weber, Economy and Society (posth., 1922)

Before a screen that repeats the same implacable phrase at any hour of the day or night — "No appointments available" — someone refreshes for the umpteenth time the immigration office's appointment page. Tomorrow, the same. Next week, the same. What that person does not see is that the appointments exist: bots capture them the instant they are published and they are resold over WhatsApp for between €50 and €350, with cases up to €600 — an entire black market built on a free formality. The bar associations, the NGOs and the Ombudsman have denounced it: the institution has urged the Interior Ministry to correct the "excessive delays" and receives constant complaints — close to 3,000 in one year, with more than 163,000 asylum applications in 2023 overwhelming the system [2]. The government has announced it will pursue the resale mafias [2]. But the underlying problem runs deeper than the mafias: when access to the administration is an opaque lottery, the citizen cannot even prove they tried.

This chapter deals with the citizen's evidentiary relationship with administrative procedure: what you filed, when you were notified, when the deadline fell. Public procurement — the administration from within — is in chapter 43; digital identity, in chapter 47. The key question: how does the citizen run a race of clocks when only one side controls the stopwatch? The answer: with a certain date of one's own at every step — the filing, the notification, the deadline, the appointment — through **t2p**: the SHA-512 fingerprint on the device, the anchoring of the fingerprint to the public blockchain and the complementary RFC 3161/eIDAS timestamp, inside an .asice container [3]. And with neutrality: the same trail that protects the citizen vouches for the diligence of the administration that does its job well and of the honest civil servant.

The architecture of deadlines

Spanish administrative procedure — Law 39/2015, of October 1 — is a machine of clocks. As a general rule, the administration has three months to decide and notify (Article 21.2), unless the procedure's specific rule sets a different term. If the term expires without a decision, administrative silence appears: the legal fiction that turns nothingness into a "yes" or a "no" depending on the procedure, with effects from the expiry; and its existence can be attested with a certificate the competent body must issue of its own motion within fifteen days — and, failing that, at the interested party's request [1].

Electronic notifications have their own clock, and it is implacable: from the moment a notification is made available in the electronic mailbox, the interested party has ten calendar days to access it; if they do not, the notification is deemed rejected and the procedure continues as if they had received it [1]. Ten days of badly chosen vacation can cost an appeal not filed, a penalty become final, a subsidy lost.

The evidentiary diagnosis: this whole edifice hangs on dates — when you filed, when it was made available, when it expires — and the dates are mostly in the custody of one of the parties. The citizen holds receipts that attest that something was filed, but not always exactly what; a mailbox whose access log belongs to the system, not to them; and a silence deadline they must compute and attest themselves. In a relationship already unbalanced in power and means, the evidentiary asymmetry is the third burden.

The appointment that is never there

The immigration case deserves its own look, because it is the paradigm of the opaque access point. The verified picture: a chronically collapsed system; bots hoarding the appointments the moment they are published; resale over WhatsApp for €50 to €350, up to €600 [2]; complaints from the bar, the NGOs and the Ombudsman, with thousands of grievances a year; more than 163,000 asylum applications in 2023, beyond the system's capacity; and the government's announcement that it will pursue the mafias [2].

But the analysis must go beyond the scandal: the black market is the symptom; the disease is the opacity of the access point. No one — not the citizen, not the Ombudsman, often not the administration itself — can audit how many appointments were published, when, and who captured them in how long. Without a verifiable record of availability, everything is suspicion: the citizen cannot prove they tried a hundred

times; the administration cannot prove it published the appointments it claims; and the bot thrives in exactly that darkness.

The structural answer is auditable availability. If every batch of published appointments and every assignment is anchored — how many, when, assigned in how long [3] — automated hoarding becomes visible to the naked eye: no human takes two hundred appointments in three seconds. The citizen can attest their dated attempts, the pursuit of the mafias [2] works on data rather than anecdotes, and reform can be demanded with figures. Transparency of the access point does not solve the underlying collapse — resources are needed [2] — but it turns a black hole into a measurable problem.

The proof of existence applied

An ordinary application, an ordinary electronic registry — and, for the first time, a file of which the citizen holds their own complete chronology.

1. The filing with sealed content

The citizen who files an application, an appeal or an allegation seals on their device the complete document and the filing receipt, together [3]: the SHA-512 fingerprint of the what and the certain date of the when. If the file ever "loses" an annex, if the version submitted is disputed, if the official receipt only attests a generic entry — the sealed copy closes the argument. And the symmetry works: the administration that anchors its own records enjoys the same protection against the citizen who claims to have filed what they never filed. No one can deny the filing; no one can invent it.

2. The clock of silence

The sealed filing starts the stopwatch with a certain date: three months — or whatever term applies [1] — counted from an indisputable fact. At expiry, the citizen holds the complete chronology to demand the silence certificate [1] or to go to court: the application existed, said this, was filed on that day, and the silence expired on this other one. Administrative silence stops being a procedural fog and becomes a dated fact — and the diligent administration, which decides within the term, holds in its own anchored records the proof of its punctuality.

3. The notification with a neutral trail

The most dangerous clock — the ten days of the electronic notification [1] — deserves a trail that depends on neither of the two parties: the availability anchored (when, exactly, it appeared in the mailbox) and the access anchored (when, exactly, it was entered) [3]. Against the "rejected for failure to access" that arrives at the worst moment, the citizen can attest what was — or was not — in their mailbox and when; against the bad-faith "it never reached me," the administration can attest the availability to the second. Defenselessness and gamesmanship die of the same medicine.

4. The auditable appointment

The anchored record of appointment availability [2][3]: every batch published, every assignment, with the exact time. The citizen who has been trying for weeks can attest their dated attempts; analysis of the record makes automated hoarding visible — the instantaneous, massive captures no human hand can perform; and the black market loses its main asset: opacity. For the administration, the same record is the defense: if it publishes what it says it publishes, the data vouch for it.

The limits, stated clearly: the seal does not resolve files or put more staff at the counter — the immigration collapse is a problem of resources and policy [2] that no fingerprint fixes; it does not replace official registries or public attestation — it complements them with a trail verifiable by both parties; and the certain date does not interpret the law: it attests facts so that lawyers, judges and the procedure itself can work on solid ground.

The promise and the peril

The promise: a citizen who runs the race of clocks with a stopwatch of their own [1][3]; access points without black holes, where the appointment resold today for €350 [2] becomes once again a free, auditable formality; and an administration that can prove its diligence with the same trail that binds it. Less

defenselessness, less gamesmanship, fewer lawsuits over facts — and more energy for the substance of the files.

The perils, with the same clarity.

The first is **the bureaucracy of the seal**. If attesting every step becomes one more burden — a formality on top of the formality — we will have worsened the relationship we wanted to heal. Sealing must be automatic and invisible: the registry that anchors on its own, the application that seals upon filing — never an added requirement that expels those with the fewest means, because the digital divide already expels enough people.

The second is **institutionalized distrust**. This chapter is not an indictment of civil servants. If the neutral trail is presented as a weapon against the administration, it will be experienced as aggression and die in a drawer; its virtue is symmetry — it protects the citizen from the lost file and the honest official from the unjust accusation. A trail of facts, never surveillance of public employees.

The third is **the mirage of the perfect counter**. No anchoring creates appointments where there are none [2] or shortens terms the law allows to be exhausted. Transparency makes the problem measurable and demandable; the solution remains political: resources, staffing, simplification.

And the fourth is **surveillance of the citizen**. A person's administrative chronology — what they apply for, what they appeal, when — sketches their life. The trails must belong to each file and its parties, never a consultable central profile; only fingerprints on the chain, content in each side's systems, and access governed by the procedure [3].

It's time to act

The most unequal relationship in the law deserves, at the very least, a shared clock.

For you, the citizen: seal what you file — document and receipt, together [3]; note and anchor the deadlines [1]; and if the system tells you there are no appointments, let your attempts be dated [2].

For lawyers and administrative agents: build the sealed chronology into every file — the filing, the notification, the silence's expiry [1][3]; it is the difference between alleging and attesting.

For the administrations: anchor your registries and notifications [3] — the neutral trail is your best defense against gamesmanship and your best credential before the citizen; and publish appointment availability in an auditable way [2].

For the Ombudsman and the oversight bodies: demand anchored records of appointment availability and assignment [2] — the 3,000 annual complaints deserve data, not form letters.

For legislators: recognize the sealed filing and notification as qualified evidence [1][3], mandate the auditability of mass access points [2] and see to it that no trail ever becomes a centralized citizen profile.

May the next person refreshing the appointments page at midnight know — and be able to prove — whether the appointment was not there or had been taken from them; and may the procedure's clock, at last, tell the same time for everyone. For the everyday counter, **time2act**.

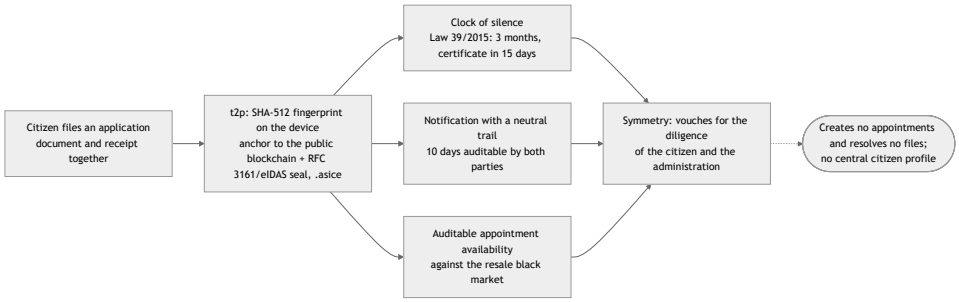


Figure 44.1 — The citizen-administration relationship with proof of existence (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, public blockchain, .asice): the filing with sealed content (what and when), the clock of administrative silence with dated deadlines (Law 39/2015: 3 months; certificate in 15 days), the notification with a neutral trail (the 10 days auditable by both parties) and auditable appointment availability (against the appointment black market). Symmetry: the same trail vouches for the administration's diligence. Limits: it creates no appointments and resolves no files; no central citizen profile. Sources: Law 39/2015; Ombudsman (immigration appointments).

[1] BOE — *Ley 39/2015, de 1 de octubre, del Procedimiento Administrativo Común de las Administraciones Públicas* (Art. 21: general three-month term to decide and notify; administrative silence with effects from expiry; silence certificate issued of the body's own motion within 15 days; electronic notifications: 10 calendar days from availability, after which deemed rejected); Administrativando — *Silencio administrativo: qué es y cuáles son sus efectos*; Conceptos Jurídicos — *Silencio administrativo en España*. https://www.boe.es/buscar/pdf/2015/BOE-A-2015-10565-consolidado.pdf · https://administrativando.es/silencio-administrativo/ · https://www.conceptosjuridicos.com/silencio-administrativo/ ↩

[2] Noticias Jurídicas — *El mercado negro de citas para extranjería: venta de citas para asilo internacional* (bots hoarding appointments; resale; system overwhelmed); OKDiario — *El caos en la web de Extranjería impulsa a las mafias que cobran hasta 350 € por cita* (WhatsApp resale from €50 to €350, up to €600); Parainmigrantes — *El problema de la cita previa llega al Defensor del Pueblo* (constant complaints, ≈3,000 in a year; >163,000 asylum applications in 2023; "excessive delays"); Iberoeconomía — *España perseguirá a las mafias que revenden citas para trámites públicos*. https://noticias.juridicas.com/contenido-patrocinado/articulo/19314-el-mercado-negro-de-citas-para-extranjeria:-venta-de-citas-para-asilo-internacional/ · https://okdiario.com/espana/caos-web-extranjeria-impulsa-mafias-que-cobran-300-cita-cada-inmigrante-13835137 · https://www.parainmigrantes.info/el-problema-de-la-gestion-de-la-cita-previa-para-tramites-de-extranjeria-llega-al-defensor-del-pueblo/ · https://iberoeconomia.es/inmigracion/espana-perseguira-a-las-mafias-que-revende-citas-para-tramites-publicos/ ↩

[3] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; content in each party's systems, only fingerprints on the chain; trail verifiable by both parties). See chapters 13 and 15; public procurement, in chapter 43. ↩

Verifiable Chance: Lotteries No One Can Load

"Babylon is nothing but an infinite game of chance."

— Jorge Luis Borges, "The Lottery in Babylon," Ficciones (1941)

A municipal hall full of people who do not usually coincide: families with folders, a notary with his own, a lottery drum — or a computer program — and a list. Subsidized housing is being allocated, and chance will decide who gets a home and who keeps waiting. It is one of the oldest and most delicate rituals of public administration. In Spain, chance decides weighty matters: the allocation of subsidized housing is done before a notary, under each administration's procedure [2]; the order of appearance of thousands of examination candidates and the composition of their tribunals come from a public draw of letters with regulated procedures — even a "mirror" system that inverts the order of the ID number so the alphabet does not always favor the same people [1]. And yet suspicion always finds a crack: in Los Llanos de Aridane (La Palma), the draw for forty homes ended up challenged, with the town council asking for it to be resolved "promptly and with maximum transparency" [2]. No fraud is needed for a draw to lose legitimacy: the reasonable doubt of whoever was unlucky is enough.

This chapter deals with public chance: how to distribute by lottery in such a way that no one can load the dice — or suspect it. Public procurement is in chapter 43; the ballot box, in chapter 48. The key question: can a draw be as verifiable as a recount? The answer: yes, with two pieces. The first, the prior commitment sealed with **t2p**: the list of participants and the allocation rule, anchored before the dice roll — SHA-512 fingerprint, RFC 3161/eIDAS timestamp, public blockchain [3]. The second, verifiable public randomness: beacons like drand, which generate random values no one can predict or bias [4]. The result: a recomputable draw — any neighbor with a computer can check that the outcome is exactly the rule applied to the list. The notary does not lose his job; he gains the best instrument in his history.

The drum and the suspicion

The traditional lottery solved the trust problem with physics: transparent drums, visible balls, a public act, cameras. Anyone can watch the drum spin, and that material display is the entire guarantee — and for decades it was enough.

But the model does not scale to administrative draws. The lists are no longer dozens of balls but thousands of names changing until the last day, managed in spreadsheets or programs only the technical staff can inspect; and the public act is often reduced to drawing a letter or a seed — everything that happens before and after is invisible [1][2]. Transparency covers the moment of the extraction; suspicion lives in everything around it.

The two weak points of any administrative draw have names. The first is **the list**: who exactly was in it when the dice rolled? A participant added late, someone unduly excluded, a duplicate — no notary can detect it if the list is a file no one has publicly fixed beforehand. The second is **the randomness**: where does the number come from? If a computer of the convening administration generates it, trust is blind; if it comes from a drum, you have to be in the room to believe it.

The Los Llanos case illustrates the cost of doubt — and precision matters: a challenge, not a proven fraud [2]. When a draw is disputed, everything stops: the families wait, the administration defends itself, and the legitimacy of the whole mechanism takes the hit. The underlying problem is not that draws are dirty; it is that they cannot prove they are clean. And in a book about proof, that distinction is everything.

The randomness no one can load

The modern solution has two halves, and the order between them is sacred.

The first half is **the prior commitment**. Before the dice roll, everything that must not be able to change is fixed: the complete list of participants with their numbers, the terms and the exact allocation rule — how one goes from the random number to the order of the list. All of this is sealed with t2p — SHA-512 fingerprint, certain date, anchoring to the public blockchain [3] — and the fingerprint is published. Consequence: no one can add, remove or reorder candidates after the fact without the fingerprint singing out; and everyone can keep the list and check that it is what it was.

The second half is **verifiable public randomness**. The random number is generated neither by the administration nor by any of its providers, but by a public beacon: drand, the randomness beacon of the League of Entropy — a coalition including Cloudflare, EPFL, Kudelski Security, Protocol Labs and the University of Chile, with later members such as the Ethereum Foundation — which emits at fixed intervals random values that are unbiased, unpredictable and verifiable by anyone, combining the entropy of multiple independent operators with threshold cryptography: neither one actor nor two in collusion can bias the result [4]. Its declared applications include, precisely, lotteries and elections [4].

And the sum is **recomputability**. It is announced in advance which future round of the beacon will be used — no one yet knows its value; the value arrives; the committed rule is applied to the committed list; and the result can be redone by anyone: with the published list, the published rule and drand's public value, every person can reproduce the entire draw on their computer and obtain exactly the same result [3] [4]. Trust stops being an act of faith in a hall: it is a check within everyone's reach.

The proof of existence applied

The same assembly hall, the same families — and a sheet with three published fingerprints anyone can check from their phone.

1. The recomputable subsidized housing

The housing draw [2] with the full protocol: the list of admitted applicants, sealed and published by fingerprint — personal data are not exposed [3]; the terms and the allocation rule, sealed; the future drand round, announced [4]; and, come the day, an allocation recomputable by anyone. In a case like Los Llanos [2], the suspicion — whichever direction it ran — could be checked in five minutes, one way or the other: either the numbers add up for everyone, or they add up for no one. The notary certifies the prior commitment and the final record: public attestation over verifiable facts.

2. Public dice for examinations

The letter for the order of appearance and the tribunal draws [1], extracted from a drand round announced in advance [4] instead of a drum in a Madrid hall: the same ritual, but verifiable from anywhere in the country by every candidate. The corrective procedures — the ID "mirror" system [1] — are published as part of the committed rule, and the whole becomes recomputable. No candidate will ever again have to wonder whether the letter came out as they say it did: they can redo the extraction themselves.

3. The sensitive assignment

The same recipe serves any random assignment that distributes valuable things or burdens: school places when a tiebreak is needed, shifts and postings, inspection samples — which companies the administration audits this year: a recomputable draw prevents both the favor and the persecution — or the order of a waiting list. Fixed rule, sealed list, public randomness [3][4]: wherever arbitrariness can dress up as chance, recomputability strips it bare. If someone comes out fifth in the audit sample, everyone can verify that that fifth place is the direct result of the calculation — and not of a phone call.

4. The record that survives

Every draw leaves its anchored record: the committed list, the rule, the random value used, the result [3]. Years later — when the appeal, the review or the retrospective suspicion arrives — the entire file can be reverified without depending on any official's memory or on the continuity of any computer system. The chance passes; its proof does not.

And the domestic version of the same problem, which fills small-claims courts every Christmas: the shared lottery ticket. The participation agreed by word of mouth — "you're in for half" — is worth whatever memory is worth when the moment comes. The photo of the ticket with the list of participants and the amounts, sealed before the draw [3], turns the country's best-loved custom into a pact with a certain date: ten seconds that save the dispute no prize deserves.

The limits, stated clearly: the seal and the beacon make the draw clean, not the policy fair — who may participate, how many homes are built or how many places are offered are human decisions no cryptography improves; recomputability requires publishing rule and list, and data protection must be respected through publication by fingerprint and pseudonymization [3]; and the notary and the public act

still make full sense — public attestation and the communal ritual are part of the legitimacy, now with mathematics underneath.

The promise and the peril

The promise: public draws anyone can recompute [3][4], challenges resolved in minutes rather than years [2], and a patrimony of legitimacy for the administration that allocates. When chance is verifiable, bad luck leaves no institutional resentment — only bad luck.

The perils, with the same clarity.

The first is **the hollowed-out ritual**. If the recomputable draw replaces the public act instead of reinforcing it, the communal dimension that legitimizes the distribution is lost — the full hall, the notary, the shared moment. The technology must accompany the ritual, not spirit it away: the act takes place, and it can additionally be checked from home.

The second is **complexity as the new opacity**. For the average citizen, "threshold cryptography" can be as impenetrable as a municipal spreadsheet. If only four experts can recompute, we will have traded one faith for another. Simple public verifiers are needed — the web page where anyone pastes their number and checks it — and pedagogy of the mechanism: theoretical verifiability is not enough; it must be practicable.

The third is **the captured seed**. Public randomness is public only if the beacon is genuinely independent and the commitment genuinely prior. Announcing the round after learning its value, or using a "beacon" controlled by the convener, would be the old fraud in a new suit [4]. The order of the steps — first the sealed commitment, then the randomness — is the whole mechanism: inverting it nullifies it.

And the fourth is **distributive techno-solutionism**. No clean draw solves the scarcity that makes it necessary. If there are forty flats for a thousand families [2], cryptography guarantees the distribution is clean — not that it is sufficient. The transparency of the mechanism must never serve to naturalize the penury of what is being distributed.

It's time to act

If chance decides lives, let it be a chance everyone can check.

For the administrations that hold draws: adopt the full protocol — list and rule sealed before, public beacon after, anchored record always [3][4]; and offer the simple verifier, so that checking does not demand engineering.

For the notaries: certify prior commitments and final records [2][3] — public attestation gains its best instrument since stamped paper.

For candidates and applicants: keep the published list and fingerprint [3]; on the day of the draw, recompute — or demand the ability to [1][4].

For the administrative courts: before a challenged draw [2], request the verifiable package — commitment, randomness, record; where it exists, the ruling is arithmetic; where it does not, the suspicion has grounds.

For civic developers: build the open verifiers — the one-click recomputation page [4]; practicable transparency is code, and not much of it.

For legislators: mandate the sealed prior commitment and verifiable public randomness in every administrative draw that distributes goods or burdens [1][2][3][4] — and make recomputability a right of every participant.

In the next hall full of families with folders, may bad luck — for it will go on existing — never again leave the aftertaste of suspicion: only the number that came up, checkable by everyone, even by those who leave without a flat. For a chance that belongs to all, **time2act**.

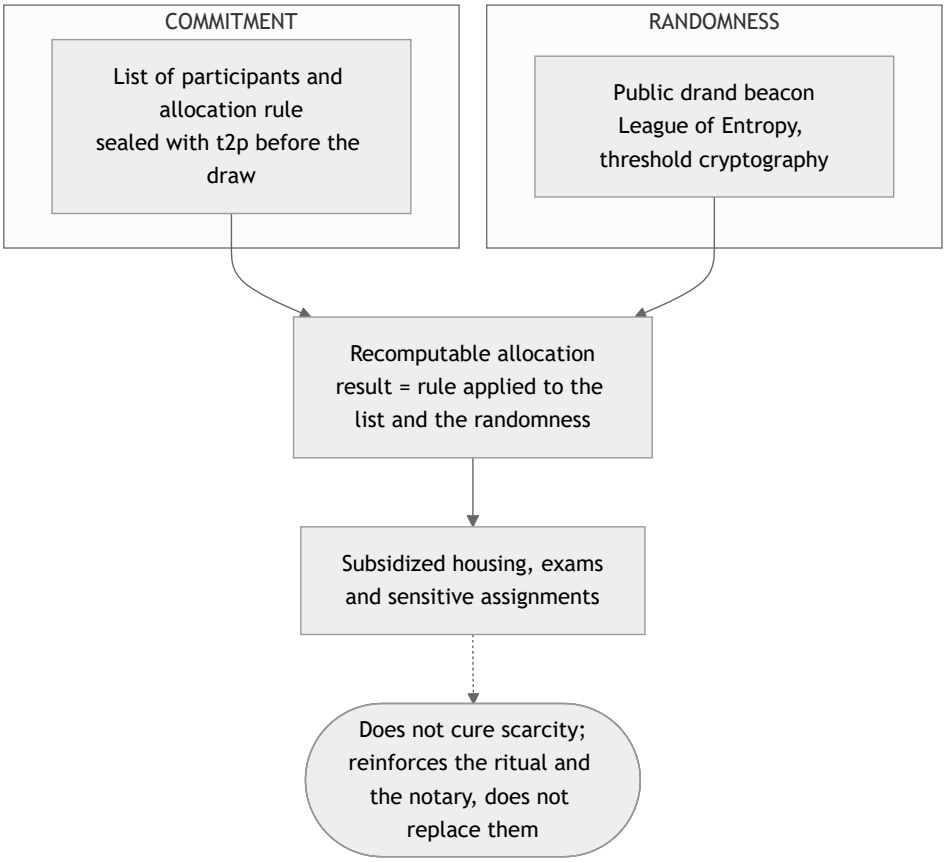


Figure 45.1 — The draw no one can load, in two halves: the prior commitment (participant list and allocation rule sealed with t2p — SHA-512 fingerprint, RFC 3161/eIDAS, public blockchain — and published by fingerprint BEFORE the draw) and verifiable public randomness (the drand beacon of the League of Entropy: unbiased, unpredictable, verifiable values generated by threshold cryptography among independent operators). Result: an allocation recomputable by anyone — result = rule(list, randomness). Limits: a clean draw does not cure scarcity; the public ritual and the notary are reinforced, not replaced. Sources: examination and subsidized-housing draws; drand/League of Entropy.

---

[1] AFAPNA — *Oposiciones 2026: sorteo de letras de tribunales y aspirantes*; OpositaTest — *Sorteo por letra del orden de actuación en oposiciones* (public letter draw; regulated subsequent procedures; "mirror" system inverting the ID order); USTEA — *Resultado del sorteo para la designación de vocales de los tribunales*. https://afapna.es/educa/oposiciones-2026-sorteo-letras-de-tribunales-y-aspirantes/ · https://blog.opositatest.com/dudas-letra-del-orden-de-actuacion-de-los-opositores/ · https://educacion.ustea.org/oposiciones-2026-resultado-del-sorteo-para-la-designacion-de-vocales-de-los-tribunales/ ↩

[2] esPúblico — *La notificación en el procedimiento de adjudicación de viviendas protegidas mediante sorteo* (draw before a notary under each administration's procedure; possibility of challenge); elDiario.es (Canarias Ahora) — *Los Llanos espera que la impugnación del sorteo de las 40 viviendas se resuelva «con prontitud y máxima transparencia»*. https://www.administracionpublica.com/adjudicaciondeviviendasprotegidas/ · https://www.eldiario.es/canariasahora/lapalmaahora/municipios/los_llanos_de_aridane/los-llanos-espera-impugnacion-sorteo-viviendas-resuelva-prontitud-maxima-transparencia_1_4583333.html ↩

[3] t2p — product manual (SHA-512 fingerprint; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; prior commitment by fingerprint — personal data are not exposed; anchored, reverifiable records). See chapters 13 and 15; electoral integrity, in chapter 48. ↩

[4] Cloudflare / League of Entropy — *Distributed Randomness Beacon* (drand: distributed, publicly verifiable randomness beacon; founding members: Cloudflare, EPFL, Kudelski Security, Protocol Labs and the University of Chile; unbiased, unpredictable, verifiable values at fixed intervals; threshold cryptography; applications: lotteries and elections); PR Newswire — *The League of Entropy launches production drand network* (the first publicly verifiable beacon in production); GitHub — *drand/drand*. https://www.cloudflare.com/leagueofentropy/ · https://www.prnewswire.com/news-releases/the-league-of-entropy-launches-production-drand-network-providing-the-first-publicly-verifiable-distributed-randomness-beacon-301109096.html · https://github.com/drand/drand ↩

Civil Service Exams: The Test Sealed Beforehand

"I agree with you that there is a natural aristocracy among men. The grounds of this are virtue and talents."

— Thomas Jefferson, letter to John Adams (28-10-1813)

A Saturday morning, in a large examination hall in any city: hundreds of people who have spent years preparing for this day, the envelope opened before everyone, the same old ritual. And the question no one can answer with proof: this exam — has anyone seen it before?

The case that makes it urgent is documented [1]. In the 2016-2017 public employment offer of Osakidetza — the Basque health service — union complaints about leaks in Anesthesia and Angiology uncovered a scandal: a report found signs of leakage in 19 exams across up to 12 medical specialties. The health minister, Jon Darpón, the director general of Osakidetza and two other executives resigned; the prosecution investigated. And the ending, eight years later: in December 2024, the judge shelved the case "as no rational indications of the commission of the offense exist" — provisionally for only two suspects from one specialty [1]. The shelving neither exonerated nor convicted: eight years of suspicion, careers destroyed, thousands of candidates left in doubt — and no way to prove anything, in either direction.

The lesson is the same as that of the aid dispute in Somalia in chapter 50: when there are no verifiable records, neither the accusation nor the defense can prove anything — and doubt devours everything, the guilty and the innocent alike.

This chapter is the natural sibling of chapter 45: there, the recomputable draw; here, the same pattern — the prior commitment — applied to the test itself. Degrees, credentials and AI in exams are in chapter 19; administrative procedure, in chapter 44; whoever reports a leak from inside, in chapter 62. The key question: how do you know the questions have not circulated, the answer key has not been touched, and your sheet is the one you handed in? The answer: sealing everything beforehand with **t2p** — the SHA-512 fingerprints of the exam and the answer key, anchored with a certain date BEFORE the sitting (only the fingerprints: the content never leaves custody), and the candidate's answer sheet, dated on the way out [2]. It protects the candidate from the dirty tribunal — and the clean tribunal from endless suspicion.

The three moments where trust breaks

The first is **before the exam**. The content is drafted, printed, kept in custody — and passes from hand to hand among drafters, reviewers and the tribunal. Every transfer is an opportunity for a copy to circulate or a version to be altered without a trace. The chain of custody is, in essence, a chain of personal trust; and when suspicion arrives — as in the Basque case [1] — there is no way to reconstruct either what existed at what moment or who had access to it.

The second is **the answer key**. It is published, the challenges arrive, questions are annulled — the legitimate routine of litigation. But it is also the moment when a key can be "adjusted": if the published version differs from the one that existed on exam day — even by just a couple of questions — the consequences move whole sets of positions. And without immutable proof of the original content, no one can show whether there was retouching or not: everything hangs on the tribunal's word.

The third is **the answer sheet**. The candidate hands in their sheet and leaves empty-handed. If the marking errs, if the sheet "gets lost," if the optical reader fails — their version exists nowhere, and the truth is orphaned.

The system asks for faith three times. And faith, as Osakidetza taught [1], once broken is not repaired by investigations: only by records that existed from the start.

The proof of existence applied

The pattern of chapter 45 — commit first, reveal later — applied to the test's content [2].

1. The exam sealed beforehand

The tribunal certifies the final exam with t2p on the day it is closed: the SHA-512 fingerprint is anchored with a certain date — and only the fingerprint: the content never leaves custody [2]. Nobody sees anything;

but from that moment, the exam handed out on day D can be checked against the prior fingerprint: no last-minute substitutions, no "versions" of unknown origin. And every access to the custody — who opened the digital envelope, when — leaves its dated record: the chain of custody stops being a chain of words.

2. The committed answer key

The marking key, sealed before the exam [2]. Challenges and annulled questions remain possible — they are the candidates' right — but they become transparent: the original key has a date; every later revision does too; and the "silent change of the correct answer" becomes impossible to perform without a trace. The tribunal that does everything right holds, for the first time, positive proof that it touched nothing.

3. The dated answer sheet

On leaving the hall, the candidate photographs their carbonless copy — or whatever receipt the system provides — with t2p: their version of the answers, with a certain date from the same day [2]. Against the marking error, the faulty optical read or the misplaced sheet, they are no longer defenseless: they hold what they handed in, dated before the key was known. And where the procedure provides no copy, this is the cheap reform to demand.

4. The challenge file — and the symmetry

When a challenge is needed, the .asice file carries everything: the prior fingerprints of the exam and the key, the dated answer sheet, the official publications captured with their dates [2] — a verifiable chronology for the administrative tribunal or the court. And the symmetry the Basque case cries out for [1]: the clean tribunal can prove the exam did not circulate and the key was not touched — the endless suspicion, which destroys innocents too, dies of dates. Osakidetza in the positive.

The limits, because they exist: the seal attests integrity and dates — not the quality of the questions or the fairness of the scoring: that belongs to the tribunal and, if need be, the judge. The prior commitment protects the content; it does not prevent the human leak — whoever has seen the exam can tell: physical custody and people still count, and the access log makes them accountable. And implementation must be institutional — the candidate can only date their own sheet: the rest is the convener's task.

The promise and the peril

The promise: civil service exams where the test's integrity is a verifiable fact rather than an act of faith [1] [2]; challenges debating the substance — whether the question was valid — over dated keys; candidates keeping their own version; clean tribunals armored against suspicion; and the next Osakidetza [1] resolved in days — in one direction or the other — instead of eight years of shadow.

And now, the perils, unsweetened.

The first is **the seal as theater**: certifying the exam while leaving physical custody as it was. The prior fingerprint is only worth something if what is handed out is genuinely checked against it, and if accesses are logged — the prior commitment is a system, not a formality.

The second is **the industrialized challenge**. If every candidate holds a file, litigation can multiply. The counterweight is the same mechanism: against dated keys and intact exams, the unfounded challenge dies of facts — those with the least reason will litigate least.

The third is **the leak no hash stops**: whoever sees the exam can memorize it and tell. The proof of existence dates and delimits responsibilities — who accessed it, when; human confidentiality is another layer, and chapter 62 explains what to do when someone, from inside, watches it break.

There is a fourth peril, quieter: that clean play makes us forget fair play. No seal makes the positions fair, the questions good or the waiting lists short. The proof of existence guarantees clean play; making the game worth playing is the task of whoever designs the calls.

It's time to act

Let the question "has anyone seen this before?" have an answer.

For tribunals and convening bodies: exam and key sealed beforehand, accesses logged, public check on day D [2] — your integrity deserves to be provable, not merely true.

For candidates: photograph and date your copy of the answers on the way out [2]; and demand the prior commitment in the calls — it is one line in the terms that changes everything.

For the unions and the academies: the cheap, cross-cutting demand: no exam without a prior fingerprint [1]; whoever refuses, let them explain why.

For the administrative courts: before a challenge, request the verifiable package — prior exam, dated key, answer sheet [2]; where it exists, the ruling is a comparison; where it does not, the suspicion has grounds.

For legislators: after Osakidetza [1] — eight years to shelve with no culprits — let the rules on access to public employment incorporate the sealed prior commitment as a requirement of every call: it costs cents and is worth the trust of millions.

In the next examination hall, when the envelope is opened, let the question "has anyone seen this before?" have an answer — a fingerprint with a date, checkable by everyone whose life is riding on it. For the positions won by studying, **time2act**.

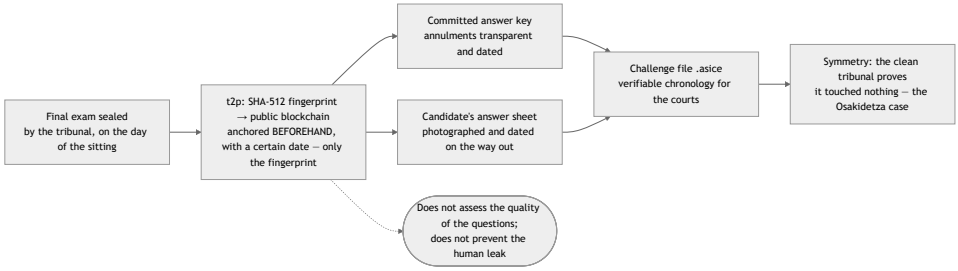


Figure 46.1 — The test sealed beforehand (t2p: SHA-512 fingerprint — only the fingerprint is anchored, the content never leaves custody — RFC 3161/eIDAS, public blockchain, .asice): the exam certified before the sitting with logged accesses, the marking key committed before the exam (annulments transparent and dated), the candidate's answer sheet dated on the way out, and the verifiable challenge file. Symmetry: the clean tribunal can prove it touched nothing — the Osakidetza case (19 exams with signs of leakage, resignations, eight years and shelving for lack of evidence) resolved in days, in either direction. Pattern: chapter 45's prior commitment, applied to the content. Sources: Redacción Médica; Deia; elDiario.es.

---

[1] Redacción Médica — *Un informe revela filtraciones de 19 exámenes en hasta 12 OPE de Osakidetza* (2016-2017 public employment offer of the Basque health service; union complaints in Anesthesia and Angiology; signs in 19 exams across 12 specialties; resignations of health minister Jon Darpón, Osakidetza's director general and two HR executives); Deia — *Archivan la causa por las presuntas filtraciones en la OPE de Osakidetza en 2018 por falta de pruebas* (December 2024: shelved "as no rational indications of the commission of the offense exist," provisional for two suspects from Plastic Surgery); elDiario.es — *La Justicia cierra sin culpables la investigación* and *Ocho años del inicio de la investigación... sin fecha para el juicio.* https://www.redaccionmedica.com/autonomias/pais-vasco/un-informe-revela-filtraciones-de-19-examenes-en-hasta-12-ope-de-osakidetza-4796 · https://www.deia.eus/actualidad/sociedad/2024/12/12/archivan-causa-filtraciones-ope-osakidetza-9042604.html · https://www.eldiario.es/euskadi/justicia-cierra-responsabilidad-investigacion-filtraciones-oposiciones-medicas-vascas_1_11896007.html ↩

[2] t2p — product manual (certify any file — only the fingerprint travels, the content never leaves custody; photo with a certain date; document scanner; capture of official publications with a date; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the tribunal draw, in chapter 45; credentials, in chapter 19; whistleblowers, in chapter 62. ↩

Digital Identity: The Proof That You Exist

"Their esse is percipi."

— George Berkeley, Principles of Human Knowledge, §3 (1710)

On September 28, 2017, in a village in the Simdega district of the Indian state of Jharkhand, an eleven-year-old Dalit girl named Santoshi Kumari died asking her mother for rice. For nearly eight days she had barely eaten. Her family's ration card — the entitlement to their quota of subsidized food — had been canceled because it was not linked to **Aadhaar**, India's biometric identity system, and for six months they had been unable to draw their grain [8]. The government disputed it, citing malaria; the right-to-food campaign documented dozens of starvation deaths tied to failures of the system [8].

We begin with this tragedy because it contains the most important warning of this chapter, and of the whole Part we now open. Identity is the right that precedes all others: without proof that you exist, there is no bank account, no vote, no school, no medical care, no aid. But a **badly designed** identity — centralized, compulsory, with no alternative — can kill. This chapter is about how to give hundreds of millions of people the proof that they exist, and how to do it **without** turning that proof into a weapon against them.

Eight hundred million invisibles

The scale of the problem is overwhelming. According to the **World Bank's Identification for Development** (ID4D) initiative, close to **800 million people** — one in ten — have no official proof of identity; the figure has fallen since 2017, but more than half of those without it are **children whose birth was never registered** [1]. **UNICEF** confirms it from the other end: **150 million children under five** are "invisible" because no one registered them at birth, and half live in sub-Saharan Africa [3]. At the other end of the spectrum lies **statelessness**: the United Nations High Commissioner for Refugees (**UNHCR**) officially counts 4.4 million at the end of 2024, but warns the real figure probably reaches ten million; the largest group are the Rohingya, nearly 1.8 million people with no recognized nationality [2].

These people do not "not exist"; **they exist**, they work, love and die, but no state recognizes it. Target 16.9 of the Sustainable Development Goals set precisely this: legal identity for all by 2030. But the key question is not *whether* we grant identity, but *how*.

Two ways of granting identity

There are two roads; the difference between them is Santoshi's lesson.

The first is the **centralized, biometric** model, exemplified by Aadhaar. With some **1.34 billion living holders** — more than 99.9% of Indian adults — and over 170 billion authentications, it is the largest identity system in history [6]. It has included many people previously excluded from the financial system. Yet in October 2023 the personal data of some **815 million Indians** — Aadhaar numbers, passports, phones, addresses — went on sale on the dark web, leaked from a covid test database [7]. Dependence on a single biometric "source of truth" excludes: authentication failure rates range between 6 and 12%, and in Rajasthan, in July 2017, one in three holders could not withdraw their grain because of repeated fingerprint failures [9]. The economist **Reetika Khera**, of the Indian Institute of Technology Delhi, puts it bluntly: the state should not treat biometric identity as the sole source of truth, because what is sold as "savings" is often the denial of rights through system failure, not fraud eliminated [9]. India's Supreme Court, which upheld Aadhaar in 2018, did so over a memorable dissent by Justice Chandrachud [9].

The second road is a new, decentralized architecture in which **the person controls their own identity**. The **W3C** standards consortium has published the technical pieces: **decentralized identifiers** (DIDs), an official recommendation since 2022, and **Verifiable Credentials**, a standard since May 2025, describing a cryptographically secure, privacy-respecting three-party ecosystem — who issues, who holds and who verifies [5]. Europe is moving in step: the **eIDAS 2.0** regulation obliges every member state to offer a **digital identity wallet** (EUDI Wallet) before the end of 2026, with **selective disclosure** built in — for example, proving you are of age without revealing your date of birth, a function already piloted in at least six states [4]. Estonia, with its e-ID covering 99% of the population, has shown

for years that it can be done well. The difference from Aadhaar is not the technology; it is **who holds the proof**.

The proof that you exist

It is on this second road that **t2p** contributes a fundamental piece: the **proof of existence of a person**. The mechanism applies the whole book's logic to the most basic case of all. The app captures a **proof of life** gesture — a video or an active challenge, resistant to deepfakes — optionally a biometric trait, and computes its **SHA-512 fingerprint** on the device itself. That fingerprint is anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS and a **forensic GPS**, and is packaged into an **.asice container verifiable offline**. The result, in plain language: *this person was alive, it was them, they were in this place, at this instant*. A signed, dated testimony for someone who has no papers.

Moreover, when an NGO, a field doctor or a school issues a **verifiable credential** — "has completed this course," "has received this vaccine" — its fingerprint can be anchored so the holder carries it **on their own device**, presents it with selective disclosure and has it verified by anyone without contacting the issuer. This is crucial for someone who loses everything in a displacement: the proof travels with the person, not in a registry that can burn or be hacked. The difference from the Aadhaar model is exactly this: **the holder keeps the proof**, not a central repository that can leak or deny access.

The promise and the peril

And here, more than in any other chapter of this book, the peril demands to be said before the promise. A proof of existence **is not a legal status**; only states grant nationality, citizenship and legal rights. t2p can prove that a person exists, who they are and where they are, but it **does not make them a citizen and does not solve statelessness**. Confusing the two would be a cruel deception.

A graver danger still, as we have seen with the 815 million leaked records and the deaths from authentication failure, is that centralizing biometrics can surveil and kill vulnerable populations. That is why the design must be radically different: **decentralized, with the proof in the holder's hands, voluntary and with a non-biometric alternative always available**. A proof of existence can never be an exclusionary requirement for accessing services; if someone declares a false name, the chain only certifies that that datum was captured intact, not its truth. The technology is good only if it serves the person, and not the state against them.

The promise is immense and deeply human. Giving 800 million invisibles a way to prove they exist — portable, theirs, verifiable without asking anyone's permission — is opening the door to all their other rights, turning "I appear in no records" into "here is the proof that I am."

An identity permanently anchored to the proof of existence can become an invisible chain preventing any personal reinvention. When the digital persona is bound "as intimately as possible" to the offline legal identity, the door opens to new injustices: life decisions — from bank loans to immigration permits — could rest on an immutable record that leaves no room for error or voluntary change [11]. The Worldcoin/World ID project illustrates this danger; its iris scanning, designed to prove humanity, has been banned or ordered to delete data in more than nine jurisdictions between 2024 and 2025 — Kenya has declared it unconstitutional and Colombia has ordered the total erasure of the biometrics collected [12]. That is why the design must include mechanisms of controlled temporary dissociation, allowing the holder to disconnect their digital identity from the legal record when they need to, and guarantee a non-biometric alternative always available, so that a single point of failure cannot block their entire future.

It's time to act

For **states and civil registries**: invest in universal birth registration and adopt decentralized models that include without excluding, learning Aadhaar's lesson. For the **humanitarian agencies** (UNHCR, NGOs): provide portable credentials controlled by the people themselves, instead of central biometric databases that turn them into subjects of surveillance. For the **standards bodies** (W3C, eIDAS): consolidate self-sovereign identity with selective disclosure as the norm. For **designers**: make privacy and voluntariness the first requirement, not the last addition. And for everyone: remember that the proof of your existence must be a right in your hands, not a chain around your neck.

A girl died because an identity system denied her the right to eat. The difference between a proof of existence that empowers and one that excludes is all the difference in the world; guaranteeing every person a secure proof, held in their own hands, is, on this fundamental front, **time2act**.

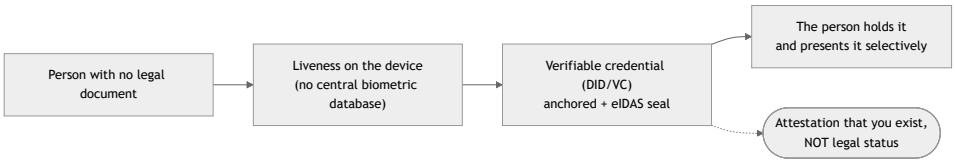


Figure 47.1 — A person without papers generates a proof-of-life gesture ("it's me, I'm alive"); on the device its SHA-512 fingerprint is computed with forensic GPS and an eIDAS timestamp ("here, now"), anchored to the public blockchain, and can accompany a W3C verifiable credential; the holder keeps the proof on their device and shares it with selective disclosure, in an .asice container verifiable offline. Scale of the challenge: ~800 M without identity, ~4.4 M (up to 10 M) stateless, 150 million unregistered children. Sources: World Bank ID4D; UNHCR; UNICEF; W3C.

---

[1] World Bank (ID4D) — *Global progress in identification: 3 findings from the latest data* (2025; ~800 M without official identity; >50% are unregistered children). https://blogs.worldbank.org/en/digital-development/global-progress-in-identification--3-findings-from-the-latest-da ↩

[2] UNHCR — *Global Trends Report 2024* (4.4 million officially stateless, up to ~10 M real; Rohingya ~1.8 M). https://www.unhcr.org/sites/default/files/2025-06/global-trends-report-2024.pdf ↩

[3] UNICEF — *Birth registration steadily increases worldwide, but 150 million children still invisible* (December 2024). https://www.unicef.org/press-releases/birth-registration-steadily-increases-worldwide-150-million-children-still-invisible ↩

[4] European Commission — *European Digital Identity (eIDAS 2.0)* (Reg. EU 2024/1183; EUDI Wallet mandatory by the end of 2026; selective disclosure). https://digital-strategy.ec.europa.eu/en/policies/eudi-regulation ↩

[5] W3C — *Verifiable Credentials Data Model v2.0 becomes a W3C Recommendation* (15-May-2025) and *Decentralized Identifiers (DID) v1.0* (2022). https://www.w3.org/press-releases/2025/verifiable-credentials-2-0/ ↩

[6] Government of India (Press Information Bureau) — *Aadhaar* (~1.34 billion living holders; >99.9% of adults; >170 billion authentications). https://www.pib.gov.in/PressReleasePage.aspx?PRID=2241778 ↩

[7] Business Standard / Resecurity — *Aadhaar data of millions of Indians put on sale on the dark web* (~815 million records, October 2023). https://www.business-standard.com/india-news/aadhaar-data-of-millions-of-indians-put-on-sale-on-the-dark-web-reports-123103000993_1.html ↩

[8] Scroll.in — *Denied food because she did not have an Aadhaar-linked ration card, Jharkhand girl dies of starvation* (Santoshi Kumari, 2017). https://scroll.in/article/854225/denied-food-because-she-did-not-have-aadhaar-linked-ration-card-jharkhand-girl-dies-of-starvation ↩

[9] Reetika Khera (IIT Delhi) — research on Aadhaar and exclusion in welfare programs (*Economic & Political Weekly*). https://www.epw.in/engage/article/aadhaar-failures-food-services-welfare ↩

[10] MIT Technology Review — *Inside the Jordan refugee camp that runs on blockchain* (WFP Building Blocks; EyePay iris recognition). https://www.technologyreview.com/2018/04/12/143410/inside-the-jordan-refugee-camp-that-runs-on-blockchain/

[11] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

[12] Biometric Update / ICLG — *Worldcoin: bans and deletion orders* (2024-2025). https://www.biometricupdate.com/202510/colombia-orders-world-shut-down-citing-biometrics-compliance-failures ↩

Electoral Integrity: Proving What Happened at the Ballot Box

"The people reign over the American political world as God rules over the universe. It is the cause and the end of all things; everything rises out of it and is absorbed back into it."

— Alexis de Tocqueville, Democracy in America, I (1835)

On September 1, 2017, in Nairobi, the Chief Justice of Kenya's Supreme Court, **David Maraga**, read out a verdict that made history. By four votes to two, the court declared the re-election of the sitting president, Uhuru Kenyatta, "invalid, null and void." It was the first time in Africa that a court annulled a president's victory at the opposition's request [1]. And the detail that brought everything down was not ballot fraud but **tally sheets that did not add up**: of 291 constituency scrutiny forms examined, 56 had no security features, 31 carried no serial number and 32 were unsigned by the agents; many results had reached the national center without the scanned images of the polling-station forms [2]. It was the integrity of the **transmission process**, not of the vote, that decided who governed Kenya.

This case marks this chapter's exact boundary. In chapter 20 we saw candidate deepfakes and electoral disinformation — the attack on the voter's *perception*. Here we speak of something else: the integrity of the **process** itself — the tally sheet, the count, the transmission, the observation — that is, being able to prove **what really happened at the ballot box**. And, as we shall see, here the most obvious technological temptation is also the most dangerous.

The process, not the ballot

The process's fragility is global. In the Democratic Republic of the Congo, in January 2019, the electoral commission proclaimed Félix Tshisekedi the winner with 38.57%. But a leaked database from the commission itself, covering 86% of the votes, gave victory to the official loser, Martin Fayulu, with nearly 60%; and 39,824 observers deployed by the Congolese Catholic Church arrived at an even higher figure for Fayulu. The two independent sources, collected separately, agreed with each other with a correlation of 0.98, yet diverged from the official result by some 3.6 million votes [3]. The evidence existed; what was missing was a verifiable chain of custody to make it incontestable.

It is precisely this chain of custody that election observers watch. As the handbook of the **OSCE/ODIHR** — the human rights office of the Organization for Security and Co-operation in Europe — explains, the value of observation lies in a systematic methodology covering the whole cycle — legal framework, registration, campaign, voting, counting, tabulation — which makes it possible to distinguish whether irregularities are **isolated or systematic** [9]. The problem is that an observer's testimony, without anchored proof, can always be questioned: "it could be from another day, another place, a setup."

Why blockchain voting is a bad idea

At this point, a warning must be issued that this book cannot spare itself, because it is counterintuitive. The natural temptation is to think: if blockchain is good for so many things, why not vote directly on it? The security community's answer is unanimous and blunt: **voting on a blockchain is a bad idea**.

The cryptographer **Ron Rivest**, of MIT and co-inventor of the RSA system, demonstrated it with his colleagues in a reference paper titled, without mincing words, *Going from bad to worse: from Internet voting to blockchain voting* (2021). Their conclusion is that internet and blockchain voting "would greatly increase the risk of **undetectable, nation-scale election failures**" [4]. The underlying reason runs deep: an election demands at once **ballot secrecy** and **software independence** — that no undetected error or change in the software can alter the result without leaving a trace — and the blockchain's transparency breaks the secrecy, while no known technology, blockchain included, guarantees software independence in mobile voting.

This is not theory. In 2018, West Virginia let 144 people — many deployed soldiers — vote with the blockchain app Voatz; in 2020, MIT researchers found that an attacker could "alter, stop or expose a user's vote," and the state abandoned it [5]. And in Moscow, in August 2019, the French researcher **Pierrick Gaudry**, of the CNRS, downloaded the code of the city's "blockchain" voting system and, in about **twenty minutes** with an ordinary computer and free software, computed its private keys: he could

have revealed every voter's ballot [6]. After the patch, another researcher broke it again. The "blockchain" label is no guarantee of security. The proof of existence, then, **must never touch the casting of the secret ballot**. Its place is elsewhere.

The proof of existence in the service of the count, not the vote

The proper place for **t2p** is the **public, observable parts** of the process. Let us see.

First, **sealing the tally sheet at the polling station**. Before transmitting any result, the official photographs the station's tally sheet and seals it: a **SHA-512 fingerprint** on the device, an **RFC 3161 timestamp** conforming to eIDAS, **forensic GPS** of the station's location and the signing official's **proof of life**, all in an **.asice** container anchored to the public blockchain. If someone alters the figure "in transit" between the station and the tabulation center — exactly Kenya's hole — the sealed original exposes it: the fingerprint does not match. Second, **the observers' evidence**: a photo or video of an irregularity sealed with GPS, time and proof of life turns a questionable testimony into a record with proven place, moment and person. Third, **the chain of custody**: every hop of the tally sheet can be re-sealed, creating the auditable trail that distinguishes the isolated from the systematic.

And here a complementarity is key. The statistician **Philip Stark**, of the University of California, Berkeley, invented **risk-limiting audits** (RLAs): a statistical sampling of the paper that formally bounds the risk of proclaiming the wrong winner, provided a **voter-verified paper ballot** exists [7]. His maxim is "don't trust the machine, demand evidence." The proof of existence does not replace the RLA: the paper remains the ground truth. t2p ensures that the *photo of the tally sheet and its transmission* are not manipulated; the authentic count is audited on the paper.

The limit must be stated with full clarity, because it is the heart of this chapter. t2p **is not internet voting and is not blockchain voting**; it only seals the public parts: tally sheets, photos, videos, transmission. And "garbage in, garbage on-chain": sealing a tally sheet proves that *that paper existed with that figure, at that time and place* — not that the figure reflects the votes actually cast. If the sheet came falsified from the station itself, the chain immortalizes it just as false. That is why the proof of existence is a **complement** to observation and to the paper audit, never their substitute.

The promise and the peril

The promise, properly bounded, is valuable. In countries where a figure altered in transit can decide who governs — and trigger violence — sealing the tally sheet at the station and giving observers incontestable evidence would make transmission fraud far harder and make it far easier to tell the isolated error from orchestrated manipulation. It is restoring credibility to a process on which peace depends.

The peril, here, is above all excess ambition. Selling the proof of existence as a solution for "voting from your phone" would repeat the mistake of Voatz and Moscow, endangering ballot secrecy and confidence in the whole system. Nor does the seal make a false tally sheet true. And any infrastructure that records electoral activity must protect and guarantee voter anonymity. The proof of existence is a tool for the transparency of what is public, not a shortcut for digitizing what must remain secret and on paper.

Furthermore, sealing tally sheets on a layer like the one t2p uses today means trusting a chain that presents itself as neutral — but neutrality is merely formal if the power to validate and interpret the protocol is concentrated in a few hands. In the L2 ecosystem, in 2026 nearly half the blocks (roughly 42%) are produced through relays that apply the OFAC lists, and all the major Ethereum L2s run on a single corporate sequencer [12]. This architecture creates an "asymmetric censorship": whoever controls the sequencer can decide which assets get recorded or excluded, without the community being aware of it. As Shoshana Zuboff reminds us, the core of power is "who knows, who decides, and who decides who decides" [11]; an anchoring infrastructure built on the public blockchain could thus replicate an oligarchic dominion, in which public trust is eroded by the absence of anti-concentration mechanisms. In that scenario, the proof of existence could end up serving more as a tool of control than as a guarantee of transparency.

It's time to act

For **electoral commissions**: seal the tally sheets at the station and transmit them, preventing the in-transit alteration that toppled Kenya. For **observers** (OSCE/ODIHR, the Carter Center, civil society): anchor your evidence with time, place and person, so that no testimony of irregularity can be dismissed as a "setup." For **auditors**: keep the voter-verified paper ballot and risk-limiting audits as the ground truth. And for everyone: resist the temptation of blockchain voting, which the best security science has advised against again and again.

A president fell in Kenya because tally sheets did not add up and no one could prove the original. Giving every tally sheet and every observation a proof with time, place and person — while never touching the secrecy of the vote — is, on this front so decisive for democracy, **time2act**.

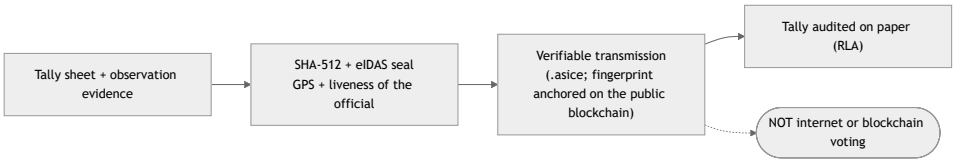


Figure 48.1 — At the station, the tally sheet is photographed and sealed with a SHA-512 fingerprint, an eIDAS timestamp, forensic GPS and the official's proof of life, in an .asice container anchored to the public blockchain; if the figure is altered in transit to the counting center, the fingerprint does not match and exposes it. What the proof of existence does NOT do: it is not internet voting, it is not blockchain voting, and it does not prove the count is true — that is audited on the paper (RLA). Sources: Supreme Court of Kenya (2017); Rivest et al. (2021); Philip Stark.

---

[1] Wikipedia (with BBC / Al Jazeera) — *2017 Kenyan presidential election petition* (4-2 annulment; CJ David Maraga; first in Africa). https://en.wikipedia.org/wiki/2017_Kenyan_presidential_election_petition ↩

[2] The Standard (Kenya) — *Presidential petition: what scrutiny of key IEBC forms revealed* (291 34B forms; 56 without security features, 31 without serial numbers, 32 unsigned). https://www.standardmedia.co.ke/article/2001253102/presidential-petition-what-scrutiny-of-key-iebc-forms-revealed ↩

[3] NYU Center on International Cooperation / Congo Research Group — *Who really won the Congolese elections?* (Fayulu ~60%; CENCO 39,824 observers; correlation 0.976-0.991; ~3.6 million votes). https://cic.nyu.edu/resources/who-really-won-the-congolese-elections/ ↩

[4] Sunoo Park, Michael Specter, Neha Narula and Ronald L. Rivest — *Going from bad to worse: from Internet voting to blockchain voting* (Journal of Cybersecurity, 2021). https://academic.oup.com/cybersecurity/article/7/1/tyaa025/6137886 ↩

[5] MIT News / USENIX Security — *Researchers find security flaws in Voatz; "The Ballot is Busted Before the Blockchain"* (2020). https://news.mit.edu/2020/voting-voatz-app-hack-issues-0213 ↩

[6] Pierrick Gaudry (CNRS/INRIA) — *Breaking the encryption scheme of the Moscow Internet voting system* (2019; broken in ~20 minutes). https://arxiv.org/abs/1908.05127 ↩

[7] Philip B. Stark (UC Berkeley) — *A Gentle Introduction to Risk-Limiting Audits* (2012). https://www.stat.berkeley.edu/~stark/Preprints/gentle12.pdf ↩

[8] Election Academy (University of Minnesota) — *Colorado successfully completes the nation's first risk-limiting audit* (2017; 9% risk limit). https://electionacademy.lib.umn.edu/2017/11/27/colorado-successfully-completes-nations-first-risk-limiting-audit/

[9] OSCE/ODIHR — *Election Observation Handbook* (6th edition; short- and long-term observers; distinguishing the isolated from the systematic). https://cdn.osce.org/sites/default/files/f/documents/5/e/68439.pdf ↩

[10] The Carter Center — *Election Standards and the Declaration of Principles for International Election Observation* (2005). https://electionstandards.cartercenter.org/our-mission/

[11] Shoshana Zuboff — *La era del capitalismo de la vigilancia* (Paidós, 2020). ↩

[12] L2BEAT; MEV Watch (2026). https://www.mevwatch.info/ ↩

Humanitarian Aid: Proving the Aid Arrives

"'Tutti fratelli' — all brothers — they kept repeating with emotion."

— Henry Dunant, A Memory of Solferino (1862)

In South Kivu, in the east of the Democratic Republic of the Congo, the organization GiveDirectly was distributing cash to families in extreme poverty through mobile phones. In 2022, part of its own local staff did something simple and devastating: they registered SIM cards in the beneficiaries' names, but kept them and put different ones in people's phones. When the payments began, the money fell onto SIMs controlled by the workers themselves. The diversion reached some **$1.2 million** and affected close to **1,900 families**; the organization confessed it publicly, but could recover only a minimal part [9]. Behind that figure is a mother who went into debt waiting for money that never came.

This case poses this chapter's question: in humanitarian aid, **how do you prove that the aid truly reaches the person who needs it?** And it also poses, with a sharpness no other chapter in the book has, its ethical reverse: the temptation to answer with biometrics can turn the most vulnerable into the most surveilled.

Ten billion that must arrive

Cash and voucher assistance — giving money instead of goods — has stopped being marginal. According to the sector's reference organization, the **CALP Network**, and Development Initiatives data, it reached some **$10.6 billion** of program spending in 2022 and already represents about a quarter of all international humanitarian aid [5][6]. The World Food Programme (WFP) is the largest provider, with $2.9 billion distributed to 51.6 million people in 76 countries in 2023 alone [5]. When that much money moves toward that many people in crisis contexts, **verifying delivery is not a detail: it is the central problem**.

The best-known technological answer is **Building Blocks**, the WFP's blockchain implementation, the largest in the humanitarian world. It began in 2017 in the Azraq camp, in Jordan, and today supports some **six million people**, with more than a million refugees using it every month [1][2]. It has processed some $555 million in transfers, eliminated 98% of banking fees in the pilot and, above all, proven its worth in **deduplication**: in Ukraine, in 2022, eighteen organizations coordinated $337 million of aid and avoided some **$35 million** of duplicated assistance by detecting overlaps between agencies [1]. Coordinating without duplicating is, in itself, a victory against fraud.

The biometric dilemma

But here comes the hard part. To know the money goes to the right person, many agencies have turned to **biometrics**. In the supermarkets of the Jordanian camps, the refugee looks at a camera and an **iris scanner** confirms their identity against the database of the High Commissioner for Refugees, which at the end of 2023 held the fingerprints and irises of **15.7 million people** [4]. The promise is impeccable: no one collects twice, no one impersonates. The problem is power. As a journalist observed in the Zaatari camp, a refugee named Bassam said the new system worked better... but he "has little say in the matter, because he has to be in the system if he wants to eat" [3]. A consent given "if you want to eat" is not free consent.

And the risk is not abstract. In Yemen, in 2019, the WFP had to partially suspend aid in Houthi-controlled areas after the Houthis refused to hand over control of the biometric data — a tool the WFP wanted precisely to stop a systematic diversion of food it had documented; the dispute still dragged on in 2023, with millions of people in need caught in the middle [8]. And the gravest case was uncovered by **Human Rights Watch**: the organization, through its crisis director **Lama Fakih**, documented in 2021 that the High Commissioner had collected biometric data from Rohingya refugees in Bangladesh and that some **830,000 names** had been shared with Myanmar — the state they had fled — without free and informed consent; 23 of 24 refugees interviewed believed their data would serve only for the aid card [7]. As the NGO **Privacy International** warns, accumulating biometrics of populations in conflict is facilitating surveillance, and when those data reach the country you are fleeing, they can be a death sentence.

Proving delivery without creating a honeypot

This is where **t2p** offers an alternative with the opposite philosophy: proving delivery **without building a central biometric database** that becomes a target.

The mechanism, at the point of delivery, is as follows. A **delivery receipt** is generated — beneficiary under an anonymous identifier, amount, place, time, agent — whose **SHA-512 fingerprint** is computed on the device, anchored to the **public blockchain** and given an **RFC 3161 timestamp** and **forensic GPS**. The result is an **.asice verifiable offline**: the proof that *this* beneficiary received *this* aid in *this* place at *this* time, auditable by the donor without having to take the implementer's word — exactly what would have exposed the Congo fraud. To confirm there is a living, consenting person, instead of a central iris scanner the **on-device proof of life** is used: the biometric template need never leave the phone or accumulate on any server.

And the most elegant piece resolves the deduplication dilemma. A team at the Swiss Federal Institute of Technology in Lausanne, led by the cryptographer **Carmela Troncoso** and in collaboration with the International Committee of the Red Cross, designed a system called **Janus** that detects double registration **without creating any cleartext biometric database**: it reveals a single bit of information — "already registered" or "new" — through secure computation [10]. It is the demonstration that phantom-beneficiary fraud can be prevented without turning refugees into a filed population. The public chain stores *proofs of uniqueness*, not identities.

And here, more than anywhere, the limit must be stated without ornament. The proof of existence demonstrates **integrity, time, place and the person's presence**, but **not the truth of what happens afterward**. If an armed actor forces the beneficiary to hand over the cash once collected — the Yemen case — or if corrupt staff control the device — the Congo case — the chain will record a "valid" delivery that was diverted: "garbage in, garbage on-chain." And there is a deeper limit still: **any biometric registry of extremely vulnerable people is surveillance**. t2p reduces the honeypot risk, but does not eliminate the dilemma of power and consent. Sometimes the safest option for the weakest is not to perfect biometrics, but to **minimize them**.

The promise and the peril

The promise is deeply human. A layer of independent, verifiable delivery proof would give donors the certainty that their money arrives, would prevent the fraud of inflated receipts and phantom beneficiaries, and — if designed with Janus's logic — would do so **without** creating the biometric databases that today endanger refugees. In a sector where every diverted dollar is a meal someone does not eat, the integrity of delivery is, literally, a matter of life.

The peril, which in this chapter walks in front, is forgetting that the technology touches the most defenseless. No system may ever be a tool of surveillance against those who flee; no consent extracted "if you want to eat" is free; and no seal stops an armed actor's bullet or the hand of a corrupt official who controls the device. The proof of existence is a tool in the service of the humanitarian principle of "do no harm," not an excuse to file the vulnerable.

Moreover, the permanent traceability of an already violated population can turn into a tradable asset: every recorded movement, every digital fingerprint, becomes a piece of "free raw material" that companies and intermediaries can package for predictive analytics markets. Shoshana Zuboff warned of this phenomenon, showing how human experience is seized and sold with no compensation to the subject [11]. If the distribution record becomes a persistent database, the same patterns that guarantee transparency could be reused for consumption profiles or even the prediction of social events, feeding a clandestine market in "vulnerability data." To prevent this commodification, explicit safeguards must be established prohibiting the sale or transfer of any record of human delivery, beyond mere on-device processing controls. Regulation must extend to the ownership and secondary use of the data, as the FTC did in 2024 by banning for the first time the sale of sensitive location data, and as the Dutch DPA did in fining Clearview AI €30.5 million over a massive biometric database [12]. This ensures the proof of existence keeps serving only the humanitarian principle of "do no harm," without opening the door to a new business model built on vulnerability.

It's time to act

For donors: demand independent, verifiable proof of delivery, so that every dollar reaches a real person and not a stolen SIM. For the agencies (WFP, UNHCR, NGOs): apply honeypot-free deduplication in the style of Janus and minimize biometrics, with genuinely free consent and non-biometric alternatives. For privacy and standards bodies: make "do no harm" and data minimization a requirement, not a recommendation. And for everyone: remember that in aid, more than anywhere, the question is not only "does it arrive?" but "does it arrive without turning the recipient into a permanent suspect?"

A Congolese mother went into debt waiting for aid that fraud stole from her. Giving every delivery an independent proof — and doing it without filing those who have already lost everything — is, on this front that touches human dignity, **time2act**.

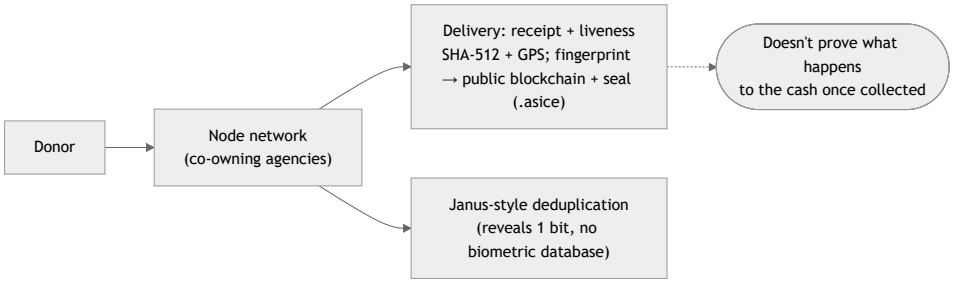


Figure 49.1 — At the point of delivery, a delivery receipt (anonymous beneficiary, amount, place, time) is sealed with a SHA-512 fingerprint, an eIDAS timestamp and forensic GPS, in an .asice container verifiable offline; the on-device proof of life confirms a living person without a central scanner; and Janus-style deduplication reveals only one bit ("already registered" or "new"), creating no biometric database. Real impact: in Ukraine ~$35 million of duplicated aid was avoided. Limit: it does not prove what happens to the cash once collected (Yemen, Congo). Sources: WFP; EPFL/ICRC (Janus); Human Rights Watch.

---

[1] WFP — *Building Blocks* (~$555 million in transfers; ~6 million people; Ukraine: $337 million coordinated, ~$35 million of overlap avoided; 159 organizations in 2025). https://www.wfp.org/building-blocks ↩

[2] UN Innovation Network — *How the Building Blocks blockchain network is transforming humanitarian aid* (>1 million refugees each month). https://www.uninnovation.network/blog/how-the-building-blocks-blockchain-network-is-transforming-humanitarian-aid ↩

[3] MIT Technology Review — *Inside the Jordan refugee camp that runs on blockchain* (EyePay; Bassam; −98% fees). https://www.technologyreview.com/2018/04/12/143410/inside-the-jordan-refugee-camp-that-runs-on-blockchain/ ↩

[4] UNHCR — *UNHCR's Biometric Tools in 2023* (database of 15.7 million people; PRIMES/BIMS). https://www.unhcr.org/blogs/unhcrs-biometric-tools-in-2023/ ↩

[5] Development Initiatives — *Cash and Voucher Assistance* (~$10.6 billion in 2022 → $9.7 billion in 2023; WFP, largest provider). https://devinit.org/resources/falling-short-humanitarian-funding-reform/cash-voucher-assistance/ ↩

[6] CALP Network — *The State of the World's Cash 2023*. https://www.calpnetwork.org/collection/the-state-of-the-worlds-cash-2023-report/ ↩

[7] Human Rights Watch — *UN Shared Rohingya Data Without Informed Consent* (Lama Fakih; ~830,000 names shared with Myanmar). https://www.hrw.org/news/2021/06/15/un-shared-rohingya-data-without-informed-consent ↩

[8] UN News — *Yemen: WFP pauses food distributions* (biometric dispute with the Houthis, 2019-2023). https://news.un.org/en/story/2023/12/1144417 ↩

[9] GiveDirectly / The New Humanitarian — *DRC mobile-cash fraud* (~$1.2 million diverted; ~1,900 families). https://www.givedirectly.org/drc-case-2023/ ↩

[10] Carmela Troncoso et al. (EPFL SPRING Lab, with the ICRC) — *Janus: Safe Biometric Deduplication for Humanitarian Aid Distribution* (IEEE S&P 2024; reveals only one bit, no biometric database). https://arxiv.org/abs/2308.02907 ↩

[11] Shoshana Zuboff — *La era del capitalismo de la vigilancia* (Paidós, 2020). ↩

[12] FTC (2024); Autoriteit Persoonsgegevens / EDPB (2024). https://www.edpb.europa.eu/news/national-news/2024/dutch-supervisory-authority-imposes-fine-clearview-because-illegal-data_en ↩

In-Kind Aid: The Chain of Custody of Relief

"To be a man is, precisely, to be responsible. [...] It is to feel, when setting one's stone, that one is contributing to the building of the world."

— Antoine de Saint-Exupéry, Wind, Sand and Stars (1939)

On June 24, 2026, at 18:04 local time, a magnitude 7.2 earthquake shakes Venezuela from near San Felipe, in the state of Yaracuy. Thirty-nine seconds later comes a second one, magnitude 7.5. Whole buildings collapse in Caracas, La Guaira and Caraballeda; Simón Bolívar International Airport suspends operations; the mobile networks saturate. As this chapter closes, the toll is overwhelming: more than 3,300 dead, more than 16,800 injured, tens of thousands missing [1]. And the world responds the way it knows how: more than 30 countries, more than 3,600 rescuers, more than 1,000 tonnes of aid. El Salvador expands its contingent from 150 to 300 rescuers and from 50 to 150 tonnes; Brazil sends 36 firefighters and six search dogs; the chef José Andrés donates a million dollars and travels to the country with World Central Kitchen; Rosalía gives two to UNICEF; Metallica, 100,000 [1][2].

And then the second question begins, the one this book cannot dodge: all of this — the tonnes, the pallets, the containers — will it reach those who need it? Because while the rescue teams are still digging through the rubble, there are already documented reports of supplies held back by police officials in Altamira, of civilian collection centers banned by mayors who restrict donations to their party's headquarters, of distributions blocked [1]. The world's generosity has begun the journey; what this chapter asks is whether the journey can be audited.

This chapter deals with in-kind aid: the things — food, medicine, tents, generators — that travel from donor to victim by sea, by air and by land. The delivery of cash aid and the deduplication of beneficiaries are in chapter 49; the journey of the donor's money — "where does my euro go" — in chapter 73; the verifiable origin of commercial products, in chapter 29; documentary fraud in international trade, in chapter 54; the real position of ships, in chapter 103; the chronology of emergencies, in chapter 105. The key question: how do you know that a pallet leaving a depot in Panama reaches a family in La Guaira intact? The answer: by having every link of the journey — the manifest, the ship or the plane, customs, the warehouse, the convoy, the distribution point — leave its proof of existence with **t2p**: the SHA-512 fingerprint computed on the device, the RFC 3161/eIDAS timestamp and the anchor to the public blockchain, inside an .asice container [11]. And with the usual neutrality: the same trail that exposes diversion protects the honest NGO and the diligent authority from false accusation. The proof serves both directions.

The precedents: when trust broke

The frauds that have marked the sector are not legends: they are documented, often by the organizations themselves. It is worth looking at them squarely, because they draw the exact map of where the chain breaks.

In November 2017, the International Federation of the Red Cross (IFRC) publicly acknowledged that fraud had taken at least $6 million of the more than $124M it had managed during the Ebola epidemic (2014-2016), the crisis that killed more than 11,000 people in West Africa. The breakdown is a lesson in dark logistics: in Sierra Leone, some $2.13 million vanished through "likely collusion" between its own staff and a local bank that set the exchange rate at the peak of the epidemic; in Guinea, some $1.17 million was lost to fraudulent invoicing by a customs-clearance provider; in Liberia, some $2.7 million evaporated in inflated prices for relief items, payrolls and volunteer incentives [5]. Honesty deserves to be said: it was the Federation itself that uncovered it and declared itself willing to waive immunities to pursue those responsible [5].

Somalia, 2010: the UN Monitoring Group reported that "up to half" of the World Food Programme's food aid was being diverted to a network of corrupt contractors, militias and local staff, and named the Adaani family — one of the program's three largest contractors — as a financier of armed groups. The WFP flatly denied it: "no evidence has been presented to us," said its executive director, Josette Sheeran, and its own audit found nothing [6]. And here lies this chapter's central evidentiary lesson: the dispute was unresolvable. Without verifiable custody records, the accusation could not prove the diversion and the defense could not prove innocence — and doubt devoured trust all the same, in both directions.

And Spain, 2007: according to the Anti-Corruption Prosecutor's Office, the executives of Intervida — one of the country's largest child-sponsorship NGOs — diverted more than €45 million disguised as "sponsorship expenses"; when the investigation became public, some 10,000 people canceled their donations. That same April, the president of Anesvad was arrested, reported by the workers of his own organization, accused of misappropriating funds [7]. The whole sector paid the trust bill.

The common pattern is what matters: the hole is never in the photo of the ship's arrival, nor in the one of the distribution with the boxes neatly stacked. It is in the dark links in the middle — the currency exchange, the customs clearance, the local contractor, the accounting between origin and destination. (The most famous case of all, the $500 million the American Red Cross raised for Haiti, belongs to the donor's side and is explained in chapter 73.)

The second catastrophe

After a disaster, aid can be lost in three documented ways — and Venezuela, by a bitter coincidence, illustrates all three.

The first is **diversion and political capture**. The verified paradigm is the CLAP program: the subsidized food boxes the Venezuelan government created in 2016, which became a vehicle for large-scale corruption. According to the investigations, the businessman Alex Saab and his partner set up a structure to buy food abroad at markups of up to $30 per box — some $360 million defrauded in two years, through companies such as Group Grand Ltd, which supplied more than 11.5 million boxes. The U.S. Treasury sanctioned the network in July 2019 for "stealing from the food distribution program"; a laboratory analysis found that the milk in the boxes had a nutritional composition below what the label claimed; and the food was used as an instrument of social control — rewarding loyalists, punishing opponents [3]. With this record, the reports of supplies held back and collection centers captured after the earthquake [1] are no surprise: they are the same logic in emergency mode.

The second is **the blockade**. In February 2019, the aid that arrived in Cúcuta, on the Colombian side of the border, found the Tienditas bridge closed off with a tanker truck and two containers welded together; on February 23, fourteen trucks carrying 280 tonnes of food tried to cross — two ended up burned in the incidents, no cargo got through, and the border was closed entirely [4]. It must be said with the neutrality this book owes itself: that shipment was also an instrument in the middle of a political dispute, and the government claimed the "aid" concealed other ends. The evidentiary lesson is not who was right; it is that, without a neutral, verifiable record of the contents of those trucks, each side had its story — and the citizen, no way of knowing. The proof of existence would have served exactly as well to show the cargo was what it claimed to be, or that it was not.

And the third is the **"second disaster."** Not everything that is lost is stolen: much is wasted. According to the humanitarian logistics expert José Holguín-Veras, between 50 and 70% of the goods that arrive in an emergency are neither needed nor appropriate; unsolicited donations clog airport runways, shelters and warehouses, and drain the volunteers who should be doing more useful things — there is a documented case of a plane that unloaded tonnes of toys onto an airport runway in the middle of an emergency, where they sat for months [8]. The honest corollary, consistent with chapter 49: the best private donation is usually money. But institutional in-kind aid — prepositioned stocks, agency cargoes — will always exist, and it is exactly the kind that needs custody.

By sea, by air, by land: the links

The real machinery of in-kind aid is more sophisticated than it looks from outside. The UNHRD network, managed by the World Food Programme, maintains prepositioned humanitarian depots in Brindisi, Accra, Kuala Lumpur, Panama City and Dubai, with medical kits, shelter material and food ready to be dispatched within 48 hours for many organizations at once; and when a crisis erupts, the Logistics Cluster led by the WFP coordinates humanitarian supply chains and fills the common gaps — shared transport, warehouses, routes [9].

Let us now follow one specific pallet. At the depot of origin it is packed and the documentation is issued: the manifest, the contents list. The pallet goes onto a ship or a plane. It arrives at a port or an airport — in the Venezuelan case, with the main airport suspended for the first days [1], every alternative route counts. It passes customs clearance. It enters a national warehouse. It leaves on trucks, often in convoy, toward

the interior. And it ends up, if all goes well, at a distribution point where someone turns it into packages for families.

The evidentiary diagnosis: every arrow of this journey is a change of hands, and every change of hands without a verifiable record is a place where things "get lost," get relabeled, get resold or get distributed by criteria other than need. Papers exist — manifests, delivery notes, receipt certificates — but they are each party's papers: modifiable, losable, contradictory. When the scandal arrives, as in Somalia [6], there is no way to reconstruct the truth; when the fraud arrives, as in Guinea's customs clearance [5], the fraudulent invoicing blends with the legitimate. What is missing is not more papers: it is papers with a certain date, a certain author and an integrity a third party can verify without asking anyone's permission.

The proof of existence applied

The same pallet, the same chain — and at every link, thirty seconds of a gesture with the phone [11].

1. The container with a pedigree

At the depot of origin, the person in charge opens t2p, turns the mode wheel to Document and scans the manifest and the contents list: the scanner detects the edges, corrects the perspective and takes up to 50 pages in a single PDF; one tap on "Certify on the blockchain" and the SHA-512 fingerprint — computed on the device, from the exact bytes of the file — receives its RFC 3161/eIDAS timestamp and is anchored to the public blockchain, with the PDF certificate and its QR code generated automatically [11]. From this instant, what the pallet contains and when it left is a fact with a certain date: the "boxes were missing" and the later relabeling are checked against the fingerprint of origin. It is the pedigree of chapter 29, applied to relief; and the sealed manifest speaks to the trade documents of chapter 54.

2. Customs and the warehouse

At every change of hands — the unloading at the port, the customs clearance, the entry into and exit from the warehouse — a photo or a scan of the certificate with t2p. The forensic GPS bar shows the location quality live: for the coordinate to be sealed at forensic grade it takes a genuine GPS fix, with an accuracy of about 50 meters and a maximum age of about 30 seconds; a simulated location is the only condition that blocks certification entirely, and if there is no signal the capture can be made anyway — the certificate honestly records that no coordinate was embedded, and why [11]. The result is an anchored chronology of entries and exits where the gaps come with dates: the pallet that went in on Monday and never came out stops being a fog and becomes a concrete question to a concrete person. Guinea's customs fraud [5] would have left exactly this trail.

3. The convoy

The trucks leave with the cargo photographed — and filmed, if necessary, with the date-and-time overlay switched on [11] — and along the route spot captures are taken with forensic GPS. The operational detail that matters in a disaster zone: t2p's certification works in the background and, if there is no coverage, the queue resumes on its own when the connection returns — you capture now, it certifies when it can, and the seal attests to the moment of capture [11]. The convoy that strays, the trailer that changes contents, the cargo that gets "redistributed" at a road checkpoint: all of it runs into a chain of anchored captures no one can rewrite. The real position of the ships carrying the big aid, by the way, is the subject of chapter 103.

4. The distribution

At the final point, the circle closes in both directions. Downward: the distribution is documented with receipts that do not expose the vulnerable — the logic of chapter 49: anonymous identifiers, no centralized biometric database. Upward: the local person in charge signs the receipt with t2p's strongest signature — the front camera records, the app poses two random liveness challenges (turn your head, smile), the device's biometric verification confirms, and "IDENTITY VERIFIED" appears [11] — verified in the sense chapter 88 makes precise: a real, present person, with the device's account and biometrics, not a name in a civil registry. And the NGO exports the .asice evidence package — a standard container (ETSI EN 319 162) with the byte-for-byte copy, the certificate, the timestamp and the verification instructions — which the donor's auditor, ten thousand kilometers away, can verify offline, without the app and without having to trust anyone [11].

The limits, with the same frankness: the seal attests to logistical facts — what, when, where, who signs — not goodness. A box of milk below its labeling [3] can be certified perfectly well: quality takes analyses and audits, not just fingerprints. No seal lifts a container welded onto a bridge [4] or disarms whoever holds back aid by force. And the trail is proportional: the cargo and the paperwork are documented, never the people who receive — here chapter 49 rules.

The promise and the peril

The promise: the donor who can follow their pallet link by link; the Somalia dispute [6], resolvable in hours against records instead of unresolvable forever; the honest NGO that answers the next scandal with evidence instead of press releases. And the sector is already walking in this direction: StaTwig traces vaccine vials from depot to health center with UNICEF's support, the IFRC has piloted blockchain with the Irish and Lebanese Red Cross societies, and the WFP is exploring food traceability from origin to destination [10]. What t2p adds is the neutral, transversal layer: the same tool for the big agency and the small NGO, with no infrastructure of their own [11].

The perils, with the same seriousness.

The first is **surveillance of the most vulnerable**. This chapter inherits, whole, the peril of chapter 49: the trail is for the cargo and the paperwork, never for turning victims into files. No centralized biometrics, no exposed beneficiary list, no condition of "if you want to eat, register." The pallet is documented, not the person who opens it.

The second is **the seal that stops no tanks**. Tienditas [4] says it all: no fingerprint dismantles a welded container. The proof of existence makes the blockade and the diversion visible and dated — and that is a lot: it turns rumor into a case file — but the answer to an armed blockade is political and diplomatic. Selling cryptography as the solution to a closed bridge would be lying.

The third is **the criminalization of the one who documents**. In hostile environments, photographing a warehouse or a road checkpoint can cost a detention, or worse. The design protects the certifier — only fingerprints on the chain, content on the device, voluntary disclosure [11] — but no protocol should force the aid worker to choose between the trail and their safety. The capture is made when it is safe to make it; the seal attests to the moment, it does not demand heroism. And when the hostile environment is outright war — when the one documenting is the journalist who is not allowed in, or the victim recording on their phone what is happening to them — this dilemma reaches its extreme: we return to it fully in chapter 113.

And the fourth is **logistical techno-solutionism**. The seal does not cure the "second disaster": if 50-70% of what arrives is not needed [8], certifying it only documents the waste with more precision. Verifiable custody works inside professional logistics — that of the depots and the cluster [9] — and with the right donation, which is often money [8]. And the final neutrality: the same trail that exposes the corrupt official exonerates the diligent one. If it only serves to accuse, it will die; if it serves to prove, it will be adopted.

It's time to act

Let generosity travel with its proof.

For donors — individuals, companies, governments: demand the verifiable chain of custody as a condition of the in-kind donation; and if you are a private individual, remember that the best help is often a monetary donation to organizations with a trail [8][11].

For NGOs and agencies: seal the manifest at origin, the certificate at every change of hands and the receipt at the final point [11]; publish by fingerprint, not by press release — you will answer the next scandal with records [5][6].

For the logistics cluster and the UN bodies: build neutral anchoring into the standards of the humanitarian chain — manifest, clearance, warehouse, distribution [9][10] — so the trail travels with the pallet between organizations.

For customs and receiving governments: emergency clearance with anchored certificates — your diligence will be dated just like any abuse [5]; and no requirement that turns the trail into a tool for

controlling those who help.

For carriers and shipping lines: cargo documented at departure and at arrival [11]; your protection against the accusation of diversion is the same trail.

For legislators and regulators of development aid: evidentiary value for anchored custody in the justification of grants; and explicit legal protection for the one who documents [7].

As this chapter is being written, pallets are still leaving for Venezuela. Let each one carry its pedigree — and let the donor, the family and the judge be able to answer the same question a year from now with the same evidence: what was there, where it went, whom it reached. So that relief arrives whole, **time2act**.

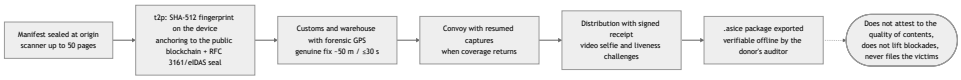


Figure 50.1 — The chain of custody of relief with proof of existence (t2p: SHA-512 fingerprint on the device, RFC 3161/eIDAS seal, public blockchain, .asice): the manifest sealed at origin (scanner of up to 50 pages), the customs and warehouse certificates with forensic GPS (genuine fix ~50 m / ≤30 s; a simulated location blocks certification), the convoy with captures resumed when coverage returns, and the receipt signed with a video selfie and liveness challenges — exportable as an .asice package verifiable offline by the donor's auditor. Limits: the seal does not attest to the quality of the contents, does not lift blockades and never files the victims. Sources: IFRC (self-disclosed Ebola fraud); UN/WFP (Somalia 2010); OFAC (CLAP network); 2026 Venezuela earthquake.

---

[1] Wikipedia (ES) — *Terremotos de Venezuela de 2026* (M7.2 San Felipe + M7.5 Yumare doublet, 24-6-2026, 39 s apart; >3,342 dead, >16,840 injured, >68,000 missing; Simón Bolívar airport suspended; distribution conflicts: supplies held back in Altamira, civilian collection centers banned in Charallave and Campo Elías, distribution blocked in El Limón, restricted access in La Guaira; response: >30 countries, >3,600 rescuers, >1,000 tonnes); ReliefWeb — *Terremoto en Venezuela: panorama de la situación (25-6-2026)*; SkyAlert — *Actualización del terremoto en Venezuela*. https://es.wikipedia.org/wiki/Terremotos_de_Venezuela_de_2026 · https://reliefweb.int/report/venezuela-bolivarian-republic/terremoto-en-venezuela-panorama-de-la-situacion-25-de-junio-de-2026 · https://skyalert.mx/articulos/actualizacion-sismos-venezuela-2026 ↩

[2] Ayuda en Acción — *Terremotos en Venezuela: qué ha pasado y situación actual* (humanitarian context; donations: José Andrés/World Central Kitchen $1 million, Rosalía $2 million to UNICEF, Metallica $100,000). https://ayudaenaccion.org/blog/ayuda-humanitaria/terremotos-venezuela/ ↩

[3] U.S. Department of the Treasury — *Treasury Disrupts Corruption Network Stealing From Venezuela's Food Distribution Program (CLAP)* (25-7-2019; Alex Saab network; Group Grand Ltd; markups); Voz de América — *¿Qué son los CLAP y qué tienen que ver con Alex Saab?*; InsightCrime — *How Criminals Take Advantage of Venezuela's Food Aid System*; Excélsior — *Despensas CLAP, una red de corrupción* (~12M boxes from Mexico, ~$30/box markup, ~$360 million; milk below labeling — Armando.Info). https://home.treasury.gov/system/files/136/2019-07-25-Venezuela-Alex-Saab-Network-SPANISH-TRANSLATION.pdf · https://www.vozdeamerica.com/a/que-son-los-clap-y-que-tienen-que-ver-con-alex-saab-/6277181.html · https://insightcrime.org/news/criminals-take-advantage-venezuelas-food-aid-system/ · https://www.excelsior.com.mx/internacional/despensas-clap-red-corrupcion-maduro-larga-historia ↩

[4] Wikipedia (ES) — *Envío de ayuda humanitaria en Venezuela de 2019* (Tienditas bridge blocked with a tanker and containers; 23-2-2019: 14 trucks/280 tonnes, 2 trucks burned; border closure); CNN en Español — *Venezuela bloquea puente en frontera con Colombia* (6-2-2019); Infobae — *Maduro ordenó retirar los contenedores del puente Tienditas* (Dec. 2022). https://es.wikipedia.org/wiki/Env%C3%ADo_de_ayuda_humanitaria_en_Venezuela_de_2019 · https://cnnespanol.cnn.com/2019/02/06/venezuela-bloquea-puente-en-frontera-con-colombia · https://www.infobae.com/america/venezuela/2022/12/14/tras-impedir-el-paso-de-la-ayuda-humanitaria-en-2019-maduro-ordeno-retirar-los-contenedores-del-puente-tienditas/ ↩

[5] NBC News — *Red Cross: $6 million for Ebola fight stolen through fraud* (Nov. 2017; ≥$6 million out of >$124 million; Sierra Leone ~$2.13M collusion with a bank over the exchange rate; Guinea ~$1.17M fraudulent invoicing by a customs-clearance provider; Liberia ~$2.7M inflated prices for relief items, payrolls and incentives); IFRC/ReliefWeb — *IFRC statement on fraud in Ebola operations*; Newsweek — *Red Cross Loses $6 Million In Fraud During Ebola Crisis*. https://www.nbcnews.com/storyline/ebola-virus-outbreak/red-cross-6-million-ebola-fight-stolen-through-fraud-n817701 · https://reliefweb.int/report/sierra-leone/ifrc-statement-fraud-ebola-operations · https://www.newsweek.com/red-cross-outraged-6-million-fraud-loss-during-ebola-outbreak-702231 ↩

[6] UN News — *UN food agency welcomes review of its food distribution programme in Somalia* (2010; Monitoring Group: "up to half" diverted; contractors/militias/local staff; Adaani family); VOA — *WFP Report says No Evidence of Food Diversion in Somalia* (WFP's denial — Josette Sheeran; own audit found no evidence); ODI — *Somalia food aid diversion*. https://news.un.org/en/story/2010/03/332792-un-food-agency-welcomes-review-its-food-distribution-programme-somalia · https://www.voanews.com/a/un-official-slams-report-on-somalia-food-diversion-89115947/153849.html · https://odi.org/en/insights/somalia-food-aid-diversion/ ↩

[7] IPS Noticias — *Corrupción-España: investigación de ONG refuerza transparencia* (2007; Intervida: >€45 million diverted according to the Anti-Corruption Prosecutor's Office, ~10,000 donors canceled; Anesvad: president arrested on 5-4-2007 for misappropriation, reported by the workers). https://ipsnoticias.net/2007/04/corrupcion-espana-investigacion-de-ong-refuerza-transparencia/ ↩

[8] Natural Hazards Center (Univ. Colorado) — *Avoiding the "Second Disaster" of Unwanted Donations* (material convergence; unsolicited donations clog runways and warehouses); TED Ideas — *After a disaster, don't send toys or clothing — send money* (José Holguín-Veras: 50-70% of the goods that arrive are neither needed nor appropriate; the toy-plane case). https://hazards.colorado.edu/news/research-counts/avoiding-the-second-disaster-of-unwanted-donations · https://ideas.ted.com/after-a-disaster-dont-send-toys-or-clothing-send-money-heres-why/ ↩

[9] WFP — *UN Humanitarian Response Depot (UNHRD)* (depots in Brindisi, Accra, Kuala Lumpur, Panama City and Dubai; dispatch within 48 h; stocks for multiple organizations); UNHRD — *Our Services*; World Economic Forum — *Logistics Emergency Team: 20 years* (Logistics Cluster led by the WFP). https://www.wfp.org/unhrd · https://unhrd.org/our-services · https://www.weforum.org/stories/global-cooperation/logistics-emergency-team-humanitarian-aid-20-years/ ↩

[10] UNICEF Office of Innovation — *Why blockchain is the future of humanitarian response* (StaTwig: traceability of vaccine vials from depot to health center); WFP Innovation Accelerator — *3 ways blockchain innovation is enhancing humanitarian work* (food traceability; Building Blocks); LSE — *How Blockchain is Contributing to the Humanitarian Sector as of 2025* (IFRC/AID:Tech pilot with the Irish and Lebanese Red Cross). https://www.unicef.org/innovation/stories/why-blockchain-future-humanitarian-response · https://wfpinnovation.medium.com/3-ways-that-blockchain-innovation-is-enhancing-humanitarian-work-e40dd3e85dee · https://blogs.lse.ac.uk/internationaldevelopment/2025/07/10/how-blockchain-is-contributing-to-the-humanitarian-sector-as-of-2025/ ↩

[11] t2p — product manual (mode wheel; document scanner up to 50 pages in one PDF; SHA-512 fingerprint computed on the device; background certification that resumes when coverage returns; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; device integrity check and hardware attestation; date/time overlay on video; signature with video selfie, liveness challenges and biometrics; .asice ETSI EN 319 162 evidence package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; cash delivery and deduplication, in chapter 49; the donor's euro, in chapter 73. ↩

Borders and Asylum: The Evidence of Those Who Flee

"We lost our home, which means the familiarity of daily life. We lost our occupation, which means the confidence that we are of some use in this world."

— Hannah Arendt, "We Refugees" (1943)

In the early hours of June 14, 2023, the fishing trawler Adriana — out of Libya five days earlier with some 750 people aboard, most from Syria, Pakistan and Egypt, children among them — capsized and sank in international waters off Pylos, Greece. More than 600 people died; 104 survived; 82 bodies were recovered [1]. The Greek coast guard denied all responsibility: the migrants had supposedly refused help, and the boat had capsized on its own. The investigations by Human Rights Watch, Amnesty International and dozens of journalists pointed in another direction: serious failures in the preceding hours and grave allegations that the Greek patrol boat caused the Adriana to capsize while trying to tow it [1][2]. Two years later, in May 2025, Greek justice indicted 17 members of the coast guard, including the patrol boat's captain, accused of causing a shipwreck resulting in the death of at least 82 people [3]. In the meantime, nine survivors had been prosecuted — facing life sentences — amid the protests of human rights organizations, who warned that the trial of the survivors could not serve as a shield against the authorities' responsibilities [2].

This chapter deals with evidence at the borders: shipwrecks, summary pushbacks and the asylum procedure. The delivery of humanitarian aid is in chapter 49; evidence of war crimes, in chapter 20; the digital identity of those without papers, in chapter 47; the real position of ships, in chapter 103. The key question: how are facts proven when one side has lost everything — often even the phone — and the other writes the official report? The answer: by leveling with proof of existence — the sealed trail of the one who flees, the anchored documentation of the observers and the equally anchored records of the authorities (**t2p**: SHA-512 fingerprint, RFC 3161/eIDAS timestamp, anchoring to the public blockchain, .asice container [5]). At the border, the neutrality of proof is not an abstract virtue: it is the only way for the truth not to depend on who tells it.

Pylos: the official version against the reconstruction

It is worth looking closely at how the truth of that night was reconstructed, because it is an entire lesson on evidence at the borders. The first version was set down by the one who had the radio and the report: according to the coast guard, the Adriana was sailing on its own and its occupants wanted no assistance [1]. Against this version there was, at the outset, nothing but 104 traumatized survivors and a boat at the bottom of the sea.

The reconstruction came through the convergence of independent sources, and it took months of work: journalists and organizations crossed the maritime traffic data of nearby vessels, the survivors' testimonies — collected one by one — and the material from their phones, to check the official chronology against that of the facts [4]. From that forensic journalism, and from the investigations of Human Rights Watch and Amnesty International, came the discrepancies the Greek investigation would eventually take up: hours without a rescue with the boat in sight, and the allegations about the fatal tow [1][2]. The procedural outcome took almost two years: on May 26, 2025, the prosecution brought charges against 17 members of the coast guard — the captain of patrol boat ΠΠΛΣ-920, accused of causing the shipwreck; all of them, of exposing people to danger despite the legal duty to rescue them — after a coalition of legal NGOs filed a complaint on behalf of 40 survivors [3].

The evidentiary lesson is the core of this chapter: in a shipwreck without anchoring, the first version is set down by whoever holds the report, the reconstruction takes years, depends on feats of journalism [4] and arrives when the front pages have moved on. If the positions, communications and images of that night had been sealed at origin — the patrol boat's and those from the passengers' phones alike [5] — the chronology would have been a fact from day one: what each party did, when, and at what distance. Chapter 103 showed how a ship's position is falsified; here we see what it costs when no one can prove theirs.

The pushback that officially never happened

The second front is summary pushbacks — the expulsions without procedure that, officially, do not exist. The work of documenting them has been done for years by civil society: the Border Violence Monitoring Network, a coalition of more than twelve organizations founded in 2017, collects testimonies of illegal pushbacks and violence at the borders of the Balkans and Greece; its "Black Book of Pushbacks" (2020) compiled the testimonies of 12,654 people, and the expanded 2022 edition runs past 3,000 pages in four volumes [6]. In these thousands of accounts, one detail repeats with chilling regularity: the confiscation or destruction of phones. The victim's proof disappears along with their only testimony.

And in December 2025 came the legal turn. The Court of Justice of the EU (case C-136/24 P, judgment of December 18, 2025) decided the case of a Syrian national who alleged that, after disembarking on a Greek island intending to apply for asylum, he had been detained, stripped of his phone and forced to re-embark out to sea until he was transferred to Turkey — with Frontex present at the operation. The Court held that the burden of proof must be adapted: demanding conclusive evidence from victims in situations of extreme vulnerability would make access to justice "illusory," and it is Frontex — responsible for the operations — that usually holds the documentation capable of clarifying the facts [7].

Read through this book's eyes, Europe's highest court has just recognized evidentiary asymmetry as a central legal problem — and confirmed that the first thing taken from the victim is, precisely, their proof. The structural answer cannot be merely procedural. If observers seal incidents in the moment [5][6], if the victim's trail exists anchored before the phone is taken [5], and if the records of official operations are equally anchored and demandable [7], the pushback "that officially never happened" becomes impossible to sustain — and the one that truly did not happen, impossible to invent. The seal has no side: that is why it can work in the most polarized territory in Europe.

Asylum: proving fear with what you have left

The third front is the quietest. Whoever applies for international protection has to substantiate persecution — the activism, the threat, the detention, the assault — with whatever they managed to carry out of a flight. Documents get lost along the way; phones are confiscated or sink [6][7]; memories contradict each other under stress; and the credibility of the account becomes the axis of the entire decision. It is, probably, the legal procedure with the widest gap between what must be proven and what can be proven.

The proof of existence applies here starting in the country of origin. Whoever fears having to flee can seal their trail before and during: the threat received, the photo of the brother's arrest, the party membership card, the published article — each element with its SHA-512 fingerprint and its certain date, anchored from the first moment [5]. Along the way, they can anchor the critical moments without publishing anything or revealing themselves.

And this is where the design shows its strength: if the phone is lost at sea or confiscated at the border [7], the content may disappear — but the existence and date of the material remain attested forever. Once a copy is recovered — from the cloud, from a relative, from an organization — its integrity and age are verified against the anchored fingerprint. The account of persecution stops depending on memory alone: it carries annexes with a certain date.

The indispensable nuance, said from the start: the seal does not make the content true. A forged document, sealed, is still forged, and the credibility assessment still belongs to the caseworker. What the seal makes impossible is retroactive fabrication — creating today "evidence" dated yesterday — which is exactly the doubt that poisons these procedures and that no interview can resolve.

The proof of existence applied

Any border, at night. Three accounts of the same moment — that of the one who flees, of the one who watches, of the one who observes — and, for the first time, none of the three can be rewritten.

1. The trail of the one who flees

The person preparing to apply for asylum seals on their device, from the country of origin, the evidence of their persecution [5]: every document, every threat, every photo, with a fingerprint and a certain date. Along the way, they can anchor the critical moments — where they were, when — without publishing anything or revealing themselves. If the phone is taken [7], the proof does not die with it: only the

fingerprint is on the chain — never the content — and it waits for the recoverable copy to attest its integrity and age. At the asylum interview, their account carries a verifiable chronology no caseworker can dismiss as "fabricated for the occasion."

2. The observer at the border

The organizations that document border incidents — the work the Black Book compiles after the fact [6] — can seal every testimony and every piece of evidence at the moment of collection [5]: the certain date of a testimony taken the day after the events is worth far more procedurally than the same testimony published two years later in a report. It is the model of chapter 20 — human rights evidence with a chain of custody — applied to the European border: fingerprint first, publication when it is safe, and an archive that the courts — which already ask Frontex for its documentation [7] — can check against the official version.

3. The night of the shipwreck, anchored

A maritime rescue with proof of existence: the positions, communications and images of all the vessels — patrol boats, merchant ships, NGO boats — sealed in real time [5], in line with the maritime integrity of chapter 103. When the dispute comes — was there a tow? when was the alert given? who was where? — the chronology will exist from the first minute: the investigation will not need two years [3], nor will the truth depend on 104 survivors being able to testify against a report [1][4]. Whoever acted well will have the best possible defense; whoever did not, no narrative escape.

4. The symmetry that protects everyone

It must be said with full clarity: the anchor also shelters the authorities. The patrol falsely accused of violence it did not commit, the officer who followed protocol, the impeccable rescue operation — all have in their own anchored records the definitive defense [5]. And the institutions gain from it what is scarcest at the borders: credibility. A coast guard with sealed logbooks does not need to be believed — it can prove it.

The limits, with the same seriousness: the seal does not decide who has the right to asylum or which migration policy is just — that belongs to laws and courts; it does not make the content of a document true, only its existence and date; it does not replace human testimony or the credibility assessment; and at the border, people's safety comes before documentation — no one should seal anything that endangers the one who flees.

The promise and the peril

The promise: borders where the facts have a chronology from the first minute [3][4], victims who do not depend on memory or on the phone that was taken from them [7], observers with unassailable archives [6] and authorities that can prove their diligence instead of merely asserting it. The CJEU has opened the legal door [7]; the infrastructure of proof is what can walk through it.

The perils, with the utmost seriousness.

The first is **surveillance of the vulnerable**. No evidence tool may become a registry of people on the move. The sealed trail belongs to the person: on their device, only fingerprints on the chain, voluntary disclosure to whomever they choose [5]. Any drift toward biometric databases of migrants would be the exact perversion of what this chapter proposes — the book has already warned of it in chapters 47 and 49.

The second is **the seal as a new border**. If asylum ended up demanding "an anchored chronology or I don't believe you," we would have raised an evidentiary wall against the most dispossessed — whoever flees a bombardment seals nothing. The proof of existence adds credibility to those who have it; its absence can never subtract it. This must be written into the rules.

The third is **the physical danger of the proof**. Documenting can kill. A phone with sealed evidence is a risk at a hostile checkpoint; protocols must put safety first — seal and delete the local content, recover it later from a safe place — and no organization should ask anyone to become an evidentiary martyr [5][6].

And the fourth is **the technological alibi**. No fingerprint rescues anyone from the sea. Rescue obligations, legal pathways and migration policy are human and political decisions [1][2][7]; the seal only

guarantees that, afterward, the truth is recoverable. Let no one sell cryptography where life vests are needed.

It's time to act

At the border, the truth cannot depend on who writes the report.

For people who fear having to flee: seal your trail — documents, threats, facts — starting now, on your device [5]; if one day your phone is taken [7], your proof will be waiting for you.

For NGOs and border observers: seal testimonies and evidence at the moment of collection [5][6]; the certain date multiplies the procedural weight of every page of your reports.

For rescue operations and coast guards: anchor logbooks, positions and communications [5] — your best defense is a chronology that cannot be disputed (chapter 103).

For Frontex and the European institutions: the CJEU already points to you as custodians of the documentation [7]; let that documentation be born anchored and be demandable — the credibility of Europe's borders depends on it.

For judges and asylum offices: weigh the sealed chronology as support for credibility [7], and never turn its absence into a presumption against.

For legislators: write the golden rule — anchored evidence adds, its absence does not subtract — and guarantee that no evidence system becomes a surveillance registry of those who flee [5].

More than 600 people died off Pylos, and it took almost two years for the investigation to point toward those who had the duty to rescue them [1][3]. Let the next dark night at sea have, from the first minute, a chronology no one can rewrite. For those who carry nothing but the truth on them, **time2act**.

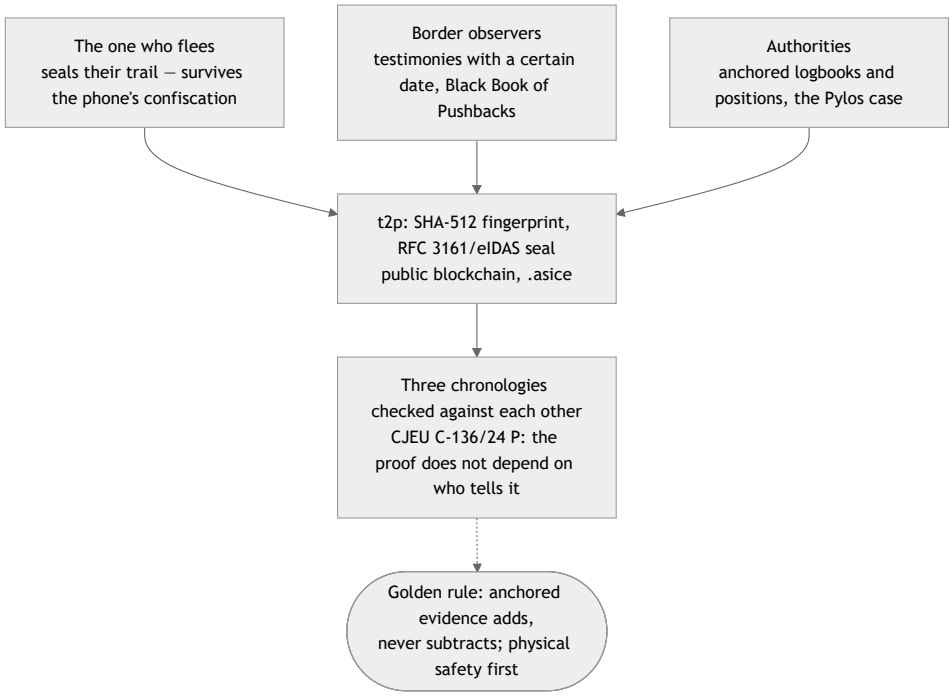


Figure 51.1 — Three anchored chronologies checked against each other (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, public blockchain, .asice): the sealed trail of the one who flees (survives the confiscation of the phone; CJEU C-136/24 P), the observers' documentation (testimonies with a certain date; Black Book of Pushbacks) and the authorities' anchored records (logbooks and positions; the Pylos case). The golden rule: anchored evidence adds credibility; its absence never subtracts. Limits: no registry of people, physical safety first, and the seal decides neither asylum nor migration policy. Sources: HRW/Amnesty (Pylos); RSA (indictment of 17 coast guard members); BVMN; CJEU C-136/24 P.

[1] Wikipedia — *2023 Pylos migrant boat disaster* (the Adriana; departure from Libya; ≈750 people aboard; more than 600 dead; 104 survivors; 82 bodies; the coast guard's version); Human Rights Watch — *Greece: Disparities in Accounts of Pylos Shipwreck Underscore the Need for Human Rights Compliant Inquiry* (the authorities' failures; allegations about the tow).

https://en.wikipedia.org/wiki/2023_Pylos_migrant_boat_disaster · https://www.hrw.org/news/2023/08/03/greece-disparities-accounts-pylos-shipwreck-underscore-need-human-rights-compliant ↩

[2] Amnesty International — *Greece: Trial of Pylos shipwreck survivors must not be a shield in the investigation of the alleged responsibilities of the Greek authorities* (May 2024; nine survivors prosecuted; criticism from human rights organizations). https://www.amnesty.org/en/latest/news/2024/05/greece-trial-of-pylos-shipwreck-survivors-must-not-be-a-shield-in-the-investigation-of-the-alleged-responsibilities-of-the-greek-authorities/ ↩

[3] Refugee Support Aegean — *Pylos Shipwreck: Criminal prosecution for felonies against 17 members of the Coast Guard* (May 26, 2025; captain of the ΠΠΛΣ-920 accused of causing a shipwreck resulting in the death of at least 82 people; complaint on behalf of 40 survivors). https://rsaegean.org/en/pylos-shipwreck-criminal-prosecution-for-felonies-to-17-members-of-the-coast-guard/ ↩

[4] Global Investigative Journalism Network — *How They Did It: Investigating the Pylos Shipwreck* (journalistic reconstruction with maritime traffic data, survivors' testimonies and material from their phones). https://gijn.org/stories/investigating-pylos-shipwreck/ ↩

[5] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; only fingerprints on the chain; optional publication; recovery of copies verifiable against the fingerprint). See chapters 13 and 15; human rights evidence, in chapter 20; the position of ships, in chapter 103. ↩

[6] Wikipedia — *Border Violence Monitoring Network* (coalition of more than 12 organizations, founded in 2017); Statewatch — *The 'Black Book of Pushbacks': testimonies of pushbacks affecting over 12,000 people* (December 2020; 12,654 people); Balkan Insight — *Study Reports 25,000 Violent Pushbacks at EU Borders* (expanded 2022 edition: more than 3,000 pages, 4 volumes). https://en.wikipedia.org/wiki/Border_Violence_Monitoring_Network · https://statewatch.org/news/2021/january/eu-the-black-book-of-pushbacks-testimonies-of-pushbacks-affecting-over-12-000-people/ · https://balkaninsight.com/2022/12/08/study-reports-25000-violent-pushbacks-at-eu-borders/ ↩

[7] Iberley — *El TJUE refuerza la tutela judicial efectiva en casos de devoluciones sumarias de migrantes* (case C-136/24 P, judgment of December 18, 2025; Syrian national stripped of his phone and returned to Turkey with Frontex present; adaptation of the burden of proof; Frontex as holder of the documentation). https://www.iberley.es/noticias/el-tjue-refuerza-tutela-judicial-efectiva-casos-devoluciones-sumarias-migrantes-35841 ↩

Wills and Inheritance: The Last Will, Made Provable

"I shall not wholly die."

— Horace, Odes, III, 30 (23 BC)

Maureen Renny, a former English school headmistress, died in January 2020 at the age of eighty-two. Shortly afterward a will appeared leaving everything to Leigh Voysey, an acquaintance of hers, and displacing an earlier will from 2016 in favor of her relatives. The document was "witnessed" by two friends of Voysey. It was false. Voysey had forged Renny's last will to claim an inheritance of **£4.2 million**, and all it took her was two complicit witnesses to fabricate it. British justice convicted her in 2024 of fraud and forgery, and also jailed her two accomplices; the criminal conviction allowed the legitimate heirs to recover the estate [1].

The case contains this chapter's singular difficulty, one no other shares: **when a will is disputed in court, the only person who could clear it up — the one who made it — is already dead**. Their voice has gone silent forever, and the document has to speak for them. This chapter is about how to make that voice, in the instant it is expressed, recorded in such a way that no one can forge it or silence it.

The will no one can defend

What is at stake is immense. The consultancy Cerulli Associates has calculated that in the United States **$84.4 trillion** will be transferred to heirs and charitable causes by 2045, in the so-called "great wealth transfer" — the largest in history [2]. And yet the foundation on which this transfer rests is surprisingly fragile. According to Caring.com's annual study, only **24% of American adults** had a will in 2025, a figure that has fallen from 33% three years earlier [3]. When there is no will, the law decides; and when there is one, the problem of proving that it is the right one begins.

As we have seen, in the United Kingdom applications to block an inheritance — *caveats* — went from 7,268 in 2019 to 11,362 in 2024, a 56% increase in five years [4]. The three classic grounds for voiding a will are always the same: lack of **testamentary capacity** (the testator was not lucid), **undue influence** (someone pressured them) and **forgery** — exactly what Voysey did [5]. Holographic wills, written by hand, are the most vulnerable, because they often carry no date and no witnesses who can testify to the testator's state at the moment of writing them.

The digital-inheritance hole

To this old problem, the digital era adds a new one. In December 2018, **Gerald Cotten**, founder of Canada's largest cryptocurrency exchange, QuadrigaCX, died at thirty during a trip to India. He alone knew the passwords to the company's cold wallets. Some **$137 million** was locked away forever, affecting more than a hundred thousand clients; his widow declared under oath that she had been unable to find the keys written down anywhere [6]. It is not an isolated case: a Chainalysis analysis estimated in 2020 that close to **20% of bitcoins** are lost and unrecoverable, often because the holder dies without sharing access [7]. As the probate and estate law section of the **American Bar Association** warns, the average person has more than a hundred online accounts but fewer than one in four has a plan for after death; without the private keys, cryptocurrency is "permanently locked, with no customer service line to call" [8]. The digital estate is, literally, buried alive.

Estonia saw it first

There is one country that understood all this before anyone else. In 2012, when the rest of the world still saw the blockchain as a bitcoin thing, Estonia put into production the planet's first state service anchored in blockchain. It was not a bank or an exchange: it was the **Succession Registry**, the registry of wills of its Ministry of Justice [9]. With the KSI technology of the company Guardtime, developed in the wake of the 2007 cyberattacks, the system seals every will by anchoring only its fingerprint — the document never leaves the state's server — so that any later alteration is detected mathematically. A small Baltic country

decided that the document most deserving of protection against manipulation is the last will of its citizens. It was right.

The proof of existence: the last will, provable

This is where **t2p** brings that intuition within reach of anyone, and it does so with a piece that shines in this chapter like in no other: the **proof of life**. At the moment of signing or recording the will, the testamentary document's **SHA-512 fingerprint** is computed on the device, anchored to the **public blockchain**, and joined by an **RFC 3161 timestamp** conforming to eIDAS, a **forensic GPS** and the **.asice container verifiable offline**. But the keystone is the **liveness recording** — a video selfie with challenges — which proves three decisive things: that the testator *was alive*, that they *were present* and that they *are who they claim to be*, at the exact instant of making the will.

Consider what this means against the two most common threats. Against **undue influence**: the testator is seen and heard declaring their will freely, with no third party directing them off-camera in any evident way. Against **forgery**: the real signer is biometrically bound to the document, so that it is no longer enough to copy a signature and produce two complicit witnesses, as in the Voysey case. And the **timestamp** resolves the eternal dispute of "a later will exists": the chain of dated fingerprints establishes an incontrovertible order of versions, revocations and codicils. Even the digital-inheritance hole has a solution: the same container can seal — encrypted — the instructions for accessing digital assets, preventing another QuadrigaCX without exposing the keys.

The limit must be stated with full honesty, because here it is subtle. The seal **does not replace** the legal formalities — witnesses, notarial intervention, probate — or validity under each jurisdiction: a will sealed but lacking the required solemnities can still be void. And, above all, **testamentary capacity is a legal and medical judgment** that the liveness video *supports* — it shows apparent lucidity, absence of visible coercion — but **does not determine**: someone can appear perfectly lucid on camera and nevertheless lack legal capacity. "Garbage in, garbage on-chain": the proof demonstrates integrity, moment, place and person, not the underlying truth of the dispositions or that no one exerted pressure off-frame.

The promise and the peril

The promise is serene and powerful. Giving every person a way to leave their last will in a form that can be proven — that it was them, alive and present, and that this is the definitive version — would save grieving families years of litigation, would protect the $84 trillion that will change hands and would close the digital-inheritance hole. It is making the testator's silenced voice keep speaking with authority.

The peril calls for nuance. The first, already said: the seal does not confer legal validity by itself and does not decide capacity; confusing the two would be dangerous. The second, the privacy of a document that is intimate by definition, which must remain in the holder's hands until the moment comes. And the third, not turning the technology into a barrier for those without access to it: the simplest way to make a will — going to the notary — must remain open to all. The proof of existence is a reinforcement of the last will, not a substitute for succession law or for the notary.

It's time to act

For individuals: make a will — most people still do not — and, when they do, seal it with proof of life and a certain date. For notaries and lawyers: incorporate the liveness recording as a reinforcement against undue influence and forgery, the most common disputes. For succession registries: follow Estonia's example and guarantee the integrity of every will. For legislators: close the legal gap of digital inheritance, so that no estate is ever again buried with a password. And for everyone: remember that the last will is the only document its author will never be able to defend in person.

A forger came close to rewriting the destiny of those £4.2 million with two complicit witnesses. Giving every testator a proof that speaks for them when they are gone — living, present and unmistakable — is, on this deeply human front, time2act.

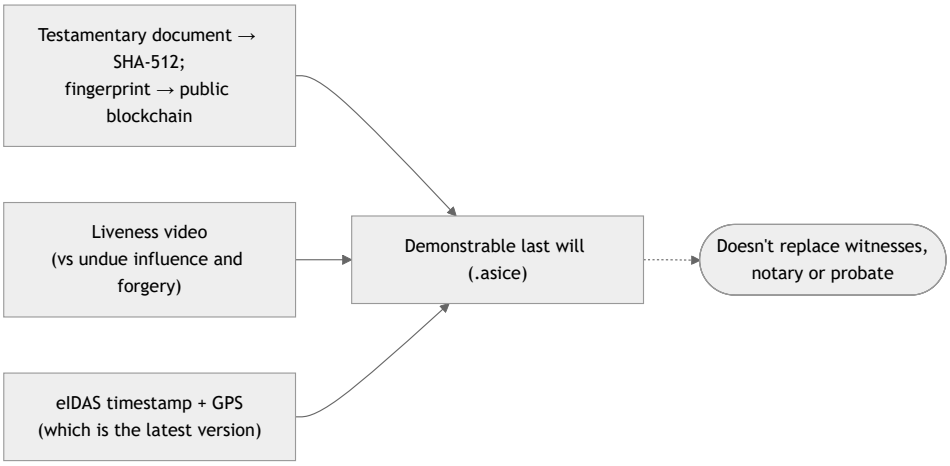


Figure 52.1 — Three inputs converge on an incontestable will: the document, whose SHA-512 fingerprint is anchored to the public blockchain (integrity); the proof of life on video, which binds identity, presence and volition (against undue influence and forgery); and the eIDAS timestamp with GPS, which fixes which version is definitive (against "a later will exists"). The result is an .asice verifiable offline. Limit: it does not replace witnesses, notary or probate, and capacity is decided by the judge or the doctor. Sources: Cerulli Associates; e-Estonia/Guardtime; the Voysey case.

---

[1] Hertfordshire Constabulary — *Three jailed for £4m inheritance fraud* (the Leigh Voysey case and Maureen Renny's will; ~£4.2 million). https://www.herts.police.uk/news/hertfordshire/news/2024/december-2024/three-jailed-for-4m-inheritance-fraud/ ↩

[2] Cerulli Associates — *Cerulli anticipates $84 trillion in wealth transfers through 2045* (2022; the "great wealth transfer"). https://www.cerulli.com/press-releases/cerulli-anticipates-84-trillion-in-wealth-transfers-through-2045 ↩

[3] Caring.com / YouGov — *Wills and Estate Planning Study 2025* (only 24% of adults have a will). https://www.caring.com/resources/wills-survey ↩

[4] Dutton Gregory / Probate Registry — *Inheritance dispute statistics* (11,362 caveats in 2024, +56% in five years). https://www.duttongregory.co.uk/site/blog/personalnews/inheritance-dispute-statistics ↩

[5] Justia — *Will contests: testamentary capacity, undue influence, forgery* (grounds for invalidity). https://www.justia.com/estate-planning/wills/will-contests/ ↩

[6] CBS News — *QuadrigaCX CEO took password to his grave* (Gerald Cotten; ~$137 million locked). https://www.cbsnews.com/news/gerald-cotten-quadrigacx-ceo-took-password-to-his-grave-leaving-company-unable-to-access-137-million/ ↩

[7] Chainalysis (via Decrypt) — *~3.7 million bitcoins lost forever* (2020; ~20% of supply). https://decrypt.co/37171/lost-bitcoin-3-7-million-bitcoin-are-probably-gone-forever ↩

[8] American Bar Association — *Digital property and estate planning* (RUFADAA; cryptocurrency is lost without the keys). https://www.americanbar.org/groups/real_property_trust_estate/resources/estate-planning/digital-property/ ↩

[9] European Commission (Interoperable Europe) / e-Estonia — *Estonia's KSI blockchain: the Succession Registry* (first state service on blockchain, 2012). https://e-estonia.com/blockchain-security-control-for-government-registers/ ↩

[10] eSignGlobal — *Electronic wills legislation and remote notarization* (eIDAS QES; wills often excluded from electronic notarization). https://www.esignglobal.com/blog/electronic-wills-legislation-remote-notary

Family Law: Co-Parenting with a Certain Date

"Your children are not your children. [...] They come through you but not from you."

— Khalil Gibran, The Prophet (1923)

A Friday afternoon, at the gate of any school. Two adults who once loved each other hand over a backpack and a child, with just enough cordiality. So far, the routine of thousands of families. The tension lies in what cannot be seen: whether the handover was at five or at half past five; whether the missing coat was with him or with her; whether the dentist's bill had been communicated; whether the weekend swap was agreed or imposed. Every minuscule detail can end up, months later, in a family court turned into "your word against mine" — with a minor in the middle.

And more and more families are in this choreography. In Spain, there were 82,991 divorces in 2024 — 8.2% more than the year before — and joint custody was granted in 49.7% of cases with children, surpassing maternal custody for the first time; in 2010 it was only 12.6% [1]. Joint custody is the new normal, and that is good news — but it multiplies the daily coordination between people who often no longer trust each other: two calendars, two homes, one parenting plan [3] and hundreds of micro-agreements a year.

This chapter deals with the evidence of co-parenting. Inheritance and the last will are in chapter 52; the judicial chain of custody, in chapter 77; the legal recording of one's own conversations, in chapter 75. The key question: how do you document a minor's life shared between two homes without turning the family into a field of mutual surveillance? The answer: with a neutral certain date on the facts both can agree to record — messages, handovers, expenses, plans — through **t2p**: the SHA-512 fingerprint on the device, the RFC 3161/eIDAS timestamp and the anchor to the public blockchain, inside an .asice container [4]. And under a principle that rules over all the others: the best interests of the child. The proof is not for winning the war; it is so there is no war.

Word against word, with a minor in the middle

The raw material of post-breakup conflicts is everyday and minuscule: handover times, extraordinary expenses, school and medical decisions, changes of plan. When the relationship is calm, none of this matters; when it deteriorates, everything becomes disputable — and the family justice system, overloaded and slow, finds itself resolving, years late, arguments about who arrived late on a Friday two years ago, with fragmentary evidence: screenshots, family testimony and memories loaded with bias.

The cost of fragile proof is paid by three parties at once. The parents, who accumulate screenshots and emails that never quite become verifiable. The court, which has to reconstruct a family's daily life from outside, deciding which fragment weighs more with no neutral layer to attest its integrity. And, above all, the minor: every hour of litigation over disputable facts is tension that reaches them, and the decisions about their life are made on versions instead of facts.

And the nuance this book cannot dodge, said with full balance: in serious conflicts — repeated breaches, crossed accusations — the lack of neutral evidence harms everyone at once. The parent who complies and cannot prove it; the one who is unjustly accused and cannot defend themselves; and the minor, whose reality is buried under the versions. The answer is not more suspicion: it is an objective record of the facts both parties have agreed to document. No more, no less.

What the courts already ask for

The judicial world has already come halfway. OurFamilyWizard — the reference co-parenting application — is used and recommended by courts and family law professionals in all 50 U.S. states and in Canada, the United Kingdom, Australia and New Zealand; its offer is exactly that of a record: messaging with a timestamp on every message and no possibility of deleting anything, read confirmations, custody calendars with a history of changes, an expense log and downloadable reports for trials and mediations — even a tone meter that flags heated language before sending [2]. In Spain, the equivalent piece on the agreement side is the parenting plan: the document that sets the common framework of upbringing beyond the divorce settlement [3].

The lesson is clear: the courts no longer want battles of memories — they want records. But today's tools have a structural dependency: the evidence lives on the provider's servers, subject to its policies, its jurisdiction and its continuity as the sole guarantee. If the company changes its terms or disappears, the record can be lost without any court holding an independently verifiable copy — chapter 75 already showed what happens when the proof depends on a contractor.

The missing layer is cryptographic neutrality: that every message, every handover confirmation and every version of the plan carry its SHA-512 fingerprint with a timestamp, anchored to the public blockchain [4] — verifiable by any court without depending on any company, or on either of the two parties. Trust stops being a hope and becomes a documented fact, and justice can devote itself to resolving real conflicts instead of questioning records.

The proof of existence applied

The same school gate, any Friday; the same two adults; and one invisible difference: the day's facts are fixed, and neither of the two will ever have to remember them again.

1. Co-parental communication with a certain date

The messages between parents — the agreements, the notices, the changes of plan — are sealed on sending: SHA-512 fingerprint, certain date, anchoring [4]. It is what the court-oriented applications already offer [2], with the neutral layer added: integrity does not depend on the provider, and the record is verifiable by any court even if the company disappears. The agreement to swap the weekend exists with a time; the notice of the medical expense, too; and the conversation that never happened cannot be invented after the fact.

2. The attested handover

The most delicate moment of the week — the handover and the pickup — can be attested with a ten-second gesture agreed by both parties: the confirmation sealed from both devices (we are here, it is this time) [4]. No camera, no recording of the other, no minor photographed: only the fact and the time, signed by both. The repeated delay gets dated; the systematic punctuality, too; and the two-year-old argument about one particular Friday disappears from the case file because the answer is already there.

3. The shared expense without a war

Extraordinary expenses — the dentist, the glasses, the summer camp — are the classic flashpoint: was it communicated? was it accepted? was it paid? With sealed receipts and notices bearing a certain date [4], the whole circuit is attested: the invoice exists, it was communicated on such a day, it was answered — or not — on such another. The parent who pays and communicates is protected; the one who answers in time, too. And silence is not a void: it is a dated record that speaks for itself, with the same neutrality for everyone.

4. The versioned parenting plan

The parenting plan [3] is a living document: it changes when the children grow, when jobs change, when life dictates. Each version sealed with a certain date [4] makes it indisputable which one was in force when: last summer's agreement, the modification agreed in September, the proposal the other party never accepted. Agreements evolve; their history does not. And when mediation or court becomes necessary, the discussion will start from the dated texts, not from opposing memories.

The limits, said with care: the seal attests to facts — times, texts, receipts — not the quality of the parenting or anyone's love; it does not decide custody — judges do that, with many more elements; it must not turn the family into a permanent dossier — what is documented is the objectifiable and agreed, never the intimacy of any home or the minor's life; and in situations of violence, the priority is immediate protection through the specialized channels, not any app.

The promise and the peril

The promise: separated families who argue less because the facts are not argued [2][4]; family courts that rule on records and not on versions; and minors who grow up with less war around them — logistics at

peace, so the adults' energy goes where it belongs.

The perils, handled with the utmost care.

The first is **the family under surveillance**. Documenting the agreed facts cannot degenerate into surveilling the other parent — not recording them, not photographing their home, not turning the minor into an informer or a carrier of tracking devices. The line is sharp: what is sealed is what both parties have agreed to document — messages, handovers, expenses, plans — and nothing else. The minor is never the object of the record; they are the reason the adults document themselves.

The second is **the procedural arsenal**. If the record is kept to stockpile ammunition, the technology will have worsened what it meant to cure. The bilateral design — both parties seal, both parties see — and the radical symmetry of the proof (it protects whoever complies, whoever that is) disincentivize its use as a weapon: against a neutral record, the best strategy is to comply.

The third is **the inequality of arms**. If only one party masters the tool, the certain date becomes a procedural advantage. As in renting and in platform work, the answer is design: free or negligible cost, simplicity, and the court weighing the bilateral record above the unilateral one.

And the fourth is **emotional techno-solutionism**. No seal cures a badly digested breakup or replaces mediation, therapy or time. The proof of existence pacifies the layer of facts; the layer of emotions is human work. Confusing them would be selling cryptography where what is needed is agreements — and, often, a little generosity.

It's time to act

Let the facts rest in peace so the people can too.

For separated parents: agree in the parenting plan [3] on what you will document — messages, handovers, expenses — and seal it bilaterally [4]; let the logistics have a certain date so the relationship does not have to keep a memory of grievances.

For family lawyers and mediators: propose the bilateral sealed record as a standard clause of the settlement [2][3]; every indisputable fact is a battle that is never fought.

For family courts: weigh bilateral anchored records above screenshots and crossed testimony [2][4]; and remind the parties that the record is a shield, not a weapon.

For co-parenting applications: build in neutral anchoring [2][4] — let your evidence not depend on your continuity as a company.

For legislators: grant qualified evidentiary value to the bilateral sealed co-parenting record, and guarantee that no tool becomes mandatory or a vehicle of surveillance inside the family [4].

At the school gate, next Friday, there will be the same two adults and the same child with their backpack. Let everything disputable be fixed in ten seconds — and let all the energy that used to burn in arguments go where it belongs: to the child watching them. For the families that go on, **time2act**.

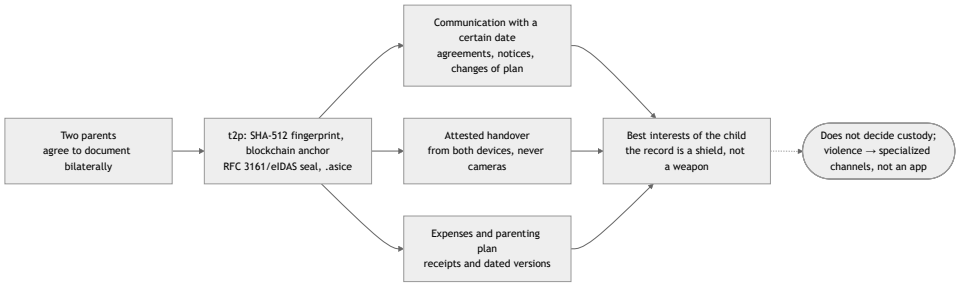


Figure 53.1 — Co-parenting logistics with proof of existence (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, public blockchain, .asice), always bilateral and agreed: communication with a certain date, the handover attested from both devices (never cameras or photographed minors), expenses with anchored receipts and the versioned parenting plan. The compass: the best interests of the child; the record is a shield, not a weapon. Limits: it does not decide custody, does not measure parenting, and in situations of violence the priority is the specialized channels. Sources: INE (joint custody 49.7%, 2024); OurFamilyWizard (courts of 50 states); the parenting plan.

[1] INE / Gracia Carrión (analysis) — *Custodia compartida supera a la materna en España (INE 2024)* (49.7% of divorces with children, +1.3 points; first time above maternal custody; 12.6% in 2010); APFS — *Separaciones y divorcios en España 2024* (82,991 divorces, +8.2%); INE — *Estadística de nulidades, separaciones y divorcios*. https://graciacarrion.com/es/custodia-compartida-supera-materna/ · https://apfs.es/separaciones-y-divorcios-en-espana-2024-tendencias-y-claves-segun-el-ine/ · https://www.ine.es/dyngs/INEbase/es/operacion.htm?c=estadistica_C&cid=1254736176798&menu=ultiDatos&idp=1254735573206 ↩

[2] OurFamilyWizard — *Tools for joint custody* (the reference co-parenting application for courts and professionals in all 50 U.S. states, Canada, the United Kingdom, Australia and New Zealand; timestamped messages with no deletion; read confirmations; calendars with history; reports for trials; ToneMeter); Crianza y Crecimiento — *Aplicaciones de coparentalidad aprobadas por tribunales*. https://es.ourfamilywizard.com/homepage · https://crianzaycrecimiento.com/aplicaciones-coparentalidad-aprobadas-tribunales/ ↩

[3] El Español — *Una abogada, sobre solicitar la custodia compartida: «Después de 10 años, necesitas un plan de parentalidad»* (the parenting plan as the common framework of upbringing beyond the divorce settlement). https://www.elespanol.com/mujer/relaciones/20251214/abogada-solicitar-custodia-compartida-despues-anos-necesitas-plan-parentalidad/1003744039064_0.html ↩

[4] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; bilateral record; content on the devices, only fingerprints on the chain; disclosure to the court). See chapters 13 and 15; the legal recording of one's own conversations, in chapter 75. ↩

Trade Finance: The Same Shipment Financed Twice

"The foundation of justice is good faith — that is, truth and fidelity to promises and agreements."

— Cicero, De Officiis, I, 23 (44 BC)

In June 2014, at the Chinese port of Qingdao, bank inspectors tried to enter the warehouses to verify something that should be elementary: that the metal listed on the invoices was actually there. They were not let in. The company Citic Resources reported that it could not locate more than one hundred thousand tonnes of alumina that should have been there [4]. It soon emerged why: the same copper and the same aluminum had been pledged, through **duplicate warehouse receipts**, to dozens of Chinese and international banks at once. The businessman behind the deception, Chen Jihong, raised some 12.3 billion yuan — close to $1.8 billion — with false documents, and ended up sentenced to twenty-three years in prison [4]. The banks' exposure exceeded $3 billion.

That scandal has a name of its own in the sector — "ghost metal" — and contains this chapter's double lesson. The first: the most expensive fraud in international trade is not faking the merchandise, but **financing the same shipment many times over** with documents no one can cross-check. The second, more uncomfortable, we will leave for the end: in Qingdao, the problem was not that the documents were false — some were perfectly real — but that the copper was not there.

The trade that lives on paper

To understand the vulnerability you have to grasp the scale. Between 80 and 90% of world trade depends on financing instruments — letters of credit, bills of lading, warehouse receipts — and the vast majority is still **paper** [10]. The International Chamber of Commerce (**ICC**), which writes the rules governing letters of credit, calculates that some **4 billion documents** circulate through the global trade system every day [10]. Yet digitalization advances at a snail's pace: even though the electronic bill of lading has existed for years, it had penetrated only **5.7%** of maritime carriers by early 2025 [3].

This dependence on paper has a brutal cost, and it is paid above all by the weakest. The **Asian Development Bank** (ADB), which takes the sector's global pulse with its reference survey, calculated that the **trade finance gap** — the difference between what companies request and what banks approve — reached a record **$2.5 trillion in 2022**, up from $1.7 trillion two years earlier [1]. Small and medium-sized enterprises, which are less than a quarter of the demand, concentrate 40% of the rejections. As the ADB sums up, financing is not lacking for lack of goods, but for lack of **verifiable trust** in the documents and the counterparty. McKinsey has estimated that digitalizing the bill of lading alone would save $6.5 billion and enable tens of billions in new trade [2].

The same shipment, financed twice

Qingdao was not an isolated case; it was the pattern. In 2020, in Singapore, the crude-trading giant Hin Leong collapsed with **$3.85 billion** of debt to twenty-three banks. The court-appointed administrators uncovered document forgery "on a massive scale": invented sales to other companies to extract almost **$112 million** from HSBC for barrels of oil that never changed hands, and it left a deep mark on the sector. Its founder, Lim Oon Kuin — "OK Lim" —, eighty-two years old, was sentenced in November 2024 to seventeen and a half years in prison; the sentence was reduced to thirteen and a half years on appeal in March 2026 [5]. And in 2021, the financing empire of **Greensill Capital** fell into insolvency, dragging down Credit Suisse, whose clients held some **$10 billion** in funds backed by invoices — some of which, according to the Swiss regulator **FINMA**, were based on transactions that had not yet occurred or were allegedly fictitious [6]. As FINMA concluded, documentary opacity, not the instrument, was what proved lethal.

In every case, the mechanism is the same: a document — a receipt, an invoice, a bill of lading — that no one can check against the other financiers, and that can therefore be presented to ten banks at once or fabricated afterward to cover a hole.

The proof of existence against double financing

This is where **t2p** strikes at the heart of the problem. At the moment of issue, the bill of lading, the commercial invoice, the warehouse receipt or the letter of credit receives a **SHA-512 fingerprint** computed on the device, is anchored to the **public blockchain**, and gains an **RFC 3161 timestamp** conforming to eIDAS and a **forensic GPS** fix of the port or warehouse location, all in an **.asice container verifiable offline**.

The blow to double financing is direct: before lending, bank B can check whether the fingerprint of that bill of lading **had already been sealed and pledged** by bank A. If the document is already "spent," the second financing is blocked — exactly what would have stopped Qingdao's duplicate receipts. The timestamp, moreover, prevents **backdating**, so that no one can fabricate afterward a document "prior" to a loss, as Hin Leong did. In this way a real invoice issued at a given moment is distinguished from one "predicted" or invented later, Greensill's problem. All of it is reinforced with the trader's **proof of life** at the moment of signing.

The limit must be stated in all its rawness: the proof of existence **does not verify that the physical merchandise exists**. In Qingdao, the document was real and could have been perfectly sealed; what was missing was the copper. "Garbage in, garbage on-chain": the chain proves the integrity, moment, place and person of the *record*, not the physical truth of the warehouse. That is why physical oracles are needed — inspections, sensors, independent custodians — which **t2p** does not replace. Furthermore, for a sealed electronic document to have enforceable "possession" and "singularity," legal recognition is required: the **United Nations Model Law on Electronic Transferable Records** (MLETR) makes it possible, but as we have seen, by 2025 only thirteen jurisdictions had adopted it, among them the United Kingdom [8]. And the lesson of the failed consortia — TradeLens, we.trade, Marco Polo, all shut down between 2022 and 2023 — shows that blocking double financing requires **all** financiers to consult the same registry. An open proof verifiable offline reduces this need for coordination, but does not eliminate it.

The promise and the peril

The promise is quantifiable: a layer of proof of existence over trade documents would make it enormously harder to finance the same shipment twice, backdate phantom sales or invent future invoices, and would help close the financing gap that, as we have seen, hits small companies in emerging countries hardest. In a sector where documentary trust is everything, sealing the integrity and moment of every document restores the foundation of trade.

The peril calls for the usual nuance. First, the seal does not make the metal appear in the warehouse; confusing the integrity of the document with the existence of the merchandise is precisely the error that cost billions in Qingdao. Second, without MLETR a sealed document is not enforceable. Third, coordination: the technology is useless if the banks do not consult it in common. The proof of existence is a necessary layer within an architecture that always includes physical inspection and the legal framework.

It's time to act

For exporters and traders: seal trade documents at origin, so that their integrity and moment are provable. For banks and financiers: check a document's fingerprint before lending, to prevent the same shipment from being financed twice. For regulators: adopt the MLETR, as the United Kingdom and a dozen jurisdictions already have, so that the sealed electronic document has full legal force. And for the sector's bodies (ICC, banks): learn from the failed consortia and converge on open standards, not closed silos.

In Qingdao, the same metal financed dozens of loans because no one could cross-check the documents. Giving every bill of lading a unique fingerprint, sealed in time and checkable by all — while never forgetting that someone has to go to the warehouse and count the copper — is, on this front of global trade, time2act.



Figure 54.1 — The exporter issues a bill of lading that is sealed with a SHA-512 fingerprint, an RFC 3161/eIDAS timestamp and GPS, anchored to the public blockchain (.asice). Bank A finances it and marks it as pledged. When a fraudster presents the "same" document to bank B, the fingerprint matches and shows it as already pledged: the second financing is blocked. Limit: the seal does not prove the merchandise physically exists in the warehouse (Qingdao's "ghost metal"). Context: $2.5 trillion trade finance gap; real electronic bills of lading, only 5.7%. Sources: Asian Development Bank; ICC; the Qingdao case.

---

[1] Asian Development Bank (ADB) — *Global trade finance gap expands to $2.5 trillion in 2022* (Trade Finance Gaps, Growth and Jobs Survey). https://www.adb.org/news/global-trade-finance-gap-expands-25-trillion-2022 ↩

[2] McKinsey — *The multi-billion-dollar paper jam: unlocking trade by digitalizing documentation* (2022; $6.5 billion in savings + new trade). https://www.mckinsey.com/industries/logistics/our-insights/the-multi-billion-dollar-paper-jam-unlocking-trade-by-digitalizing-documentation ↩

[3] DCSA / FIT Alliance — *eBL adoption* (real penetration of the electronic bill of lading: 5.7%, January 2025). https://dcsa.org/standards/booking/adoption-guide-booking-and-bill-of-lading ↩

[4] Caixin Global — *Dezheng founder jailed over Qingdao port forgery scandal* (Chen Jihong, 23 years; ~12,300M yuan with false receipts). https://www.caixinglobal.com/2018-12-12/dezheng-founder-gets-32-years-in-prison-over-port-forgery-scandal-101358154.html ↩

[5] Singapore Police Force / SCMP — *Hin Leong founder OK Lim jailed* (17.5 years, reduced to 13.5 on appeal in March 2026; ~$112 million from HSBC). https://www.police.gov.sg/media-hub/news/2024/20241118_founder_of_hin_leong_trading_sentenced_and_jailed_for_cheating_and_instigating_forgery ↩

[6] FINMA — *Greensill / Credit Suisse* (2021 insolvency; ~$10 billion in four funds; fictitious invoices). https://www.finma.ch/en/news/2023/02/20230228-mm-greensill/ ↩

[7] Maersk — *Maersk and IBM to discontinue TradeLens* (consortium shutdown, 2022-2023). https://www.maersk.com/news/articles/2022/11/29/maersk-and-ibm-to-discontinue-tradelens

[8] UNCITRAL — *Model Law on Electronic Transferable Records (MLETR): status* (13 jurisdictions; including the United Kingdom). https://uncitral.un.org/en/texts/ecommerce/modellaw/electronic_transferable_records/status ↩

[9] Clifford Chance — *UK Electronic Trade Documents Act 2023* (legal equivalence of the electronic bill of lading). https://www.cliffordchance.com/briefings/2023/08/uk-electronic-trade-documents-act-2023--a-further-step-towards-p.html

[10] International Chamber of Commerce (ICC) — *Survey shows steady rise of global adoption of electronic bills of lading* (~4 billion documents/day; UCP 600 rules). https://iccwbo.org/news-publications/news/survey-shows-steady-rise-of-global-adoption-of-electronic-bills-of-lading/ ↩

Audit: Accounts and Reports Proof Against Tampering

"What advantages does he derive from the system of book-keeping by double entry! It is among the finest inventions of the human mind."

— Goethe, Wilhelm Meister's Apprenticeship (1796)

There was a fiduciary account that did not exist. For years, the German payments company Wirecard declared that it held **€1.9 billion** — almost a quarter of its balance sheet — deposited in escrow accounts at banks in the Philippines. When EY's auditors asked the Singapore trustee, Shan Rajaratnam, he told them the truth: he held no Wirecard money in custody. But some Wirecard managers intervened, and Shan "corrected" his answer and confirmed the 1.9 billion. EY settled for screenshots and documents supplied by the client itself, and **never wrote directly to the banks** [1]. Two years later, the banks confirmed that those statements were forged and that the account had never existed. Wirecard went bankrupt in June 2020, in the biggest fraud of postwar Germany.

The detail that makes it so painful is its simplicity: a **direct confirmation letter from the bank** — what a rival auditor described as "day-one training at audit school" — would have been enough to bring the whole deception down. This chapter is about exactly that: how to make the evidence on which trust in a company's accounts rests impossible to fabricate.

The fraud the auditors did not see

Wirecard is not an exception; it is the latest link in a sadly famous chain. In 2001, Enron collapsed and evaporated more than **$74 billion** of its shareholders' money; its auditor, Arthur Andersen — one of the world's five largest, with more than **85,000 employees** — was found guilty of destroying documents and disappeared, giving rise to the Sarbanes-Oxley Act [4]. In 2020, the coffee chain Luckin Coffee admitted to having fabricated some **$311 million** in sales [5], and the Indian case of Satyam had already shown Wirecard's exact pattern: false invoices and **forged bank confirmations**. According to the Association of Certified Fraud Examiners, organizations lose on average 5% of their revenue to fraud, and financial statement fraud, despite being only 5% of the cases, is the most expensive, with a median loss of **$766,000** per case [7].

The common thread of all these scandals is that the auditor works with documents that have passed through the auditee's hands; when the auditee is the fraudster, the evidence arrives already contaminated. It is worth remembering who uncovered Wirecard: the journalist **Dan McCrum**, of the *Financial Times*, who in his book *Money Men* (2022) recounts how he was surveilled, hacked and even criminally investigated for "market manipulation" — **for being right** — while the company kept trading on the German index [3]. The system punished the messenger before the fraud.

The regulator responds, and the new ESG front

The regulatory response has moved, precisely, toward the evidence. The American audit supervisor, the **PCAOB**, adopted in 2023 a specific standard — AS 2310, "The Auditor's Use of Confirmation" — reinforcing that confirmations must come directly from reliable external sources and be evaluated with skepticism [6]. It is, at bottom, an acknowledgment that the direct bank confirmation is the link that fights corporate fraud.

And a new, even more fragile front is opening: sustainability. The European **CSRD** directive was born to oblige tens of thousands of companies to file a sustainability report with **external assurance** — independent verification — under the ESRS standards developed by **EFRAG** [8]. In 2025-26, the "Omnibus" package cut its scope by 80%, limiting it to companies with more than a thousand employees and €450 million in turnover, but kept limited assurance mandatory [9]. And the world audit standards board, the **IAASB**, approved in 2024 the first global standard for verifying sustainability information, ISSA 5000 [10]. The underlying problem is that ESG data — emissions, supply chain data — is the terrain where **data provenance** is the weakest of all. Verifying it requires knowing where it comes from.

The proof of existence: sealing at the source

This is where **t2p** offers the natural technical complement to all this regulation. The idea is to apply the seal **at the point of capture**, not once the document has already passed through the auditee. A bank confirmation signed and sealed **at the bank itself** — with a SHA-512 fingerprint, anchoring to the public blockchain and an RFC 3161 timestamp conforming to eIDAS, in an .asice container verifiable offline — cannot be replaced by a fabricated screenshot or redirected to a complicit trustee. The auditor stops trusting an intermediated piece of paper and verifies an **unalterable chain of custody** from the source.

The same applies to the accounting books, to invoices, to physical inventory counts — with GPS and the proof of life of whoever performs them — and, above all, to the **auditor's working papers**, sealed at the moment of creation. Against the Andersen case, this makes it **impossible to destroy them or backdate their content** without the anchor giving it away. And for the CSRD's most fragile front, sealing the ESG datum at the source — the emissions sensor reading, the supplier's invoice — is the only way to give solid provenance to what today is verified almost blind.

The limit must be said in all its rawness, because in auditing it is essential. The proof of existence demonstrates the **integrity, moment, place and person** of a record, not the underlying **economic truth**. If all the inputs are fabricated upstream — a complicit bank sealing a false statement from the source, invented sales as at Luckin, a special-purpose entity as at Enron — the seal will only certify that **a lie was captured unalterably**: "garbage in, garbage on-chain." t2p drastically reduces subsequent falsification and eliminates the alibi of the "lost or altered document," but it **does not replace the auditor's judgment or confirmation with a genuinely independent external source**. A seal does not turn a false datum into a true one.

The promise and the peril

The promise is solid and quantifiable. A layer of proof of existence over confirmations, books and working papers would make financial statement fraud — which costs billions and destroys jobs and savings — much harder; it would give the auditor a reliable chain of custody instead of intermediated documents, and it would lay a solid foundation for sustainability assurance, where data provenance is weakest. In a system built on trust in the accounts, making the evidence unalterable reinforces the whole edifice.

The peril calls for the usual nuance. The first, as we have said: the seal does not turn a false datum into a true one; if the source is already complicit, the chain immortalizes the lie. Confusing documentary integrity with economic truth would be precisely Wirecard's fundamental error. The second, overconfidence: no technology replaces the auditor's professional skepticism. The proof of existence is a tool that reinforces external confirmation and protects the working papers, not a substitute for human judgment or for the independence of the source.

It's time to act

For auditors: demand confirmations sealed at the source and seal your own working papers as you create them, so that no one can fabricate or destroy them. For banks and external sources: issue confirmations with a cryptographic seal from the source, not through the client. For regulators (PCAOB, IAASB, EFRAG): recognize the seal of origin as a reinforcement of AS 2310 and of sustainability assurance. And for companies: seal the ESG datum at the source, before any auditor has to verify it blind.

An account that did not exist nearly went unnoticed because no one wrote to the bank. Giving every confirmation, every book and every working paper a proof sealed at the source — and always remembering the need for professional skepticism — is, on this front that sustains trust in the markets, **time2act**.

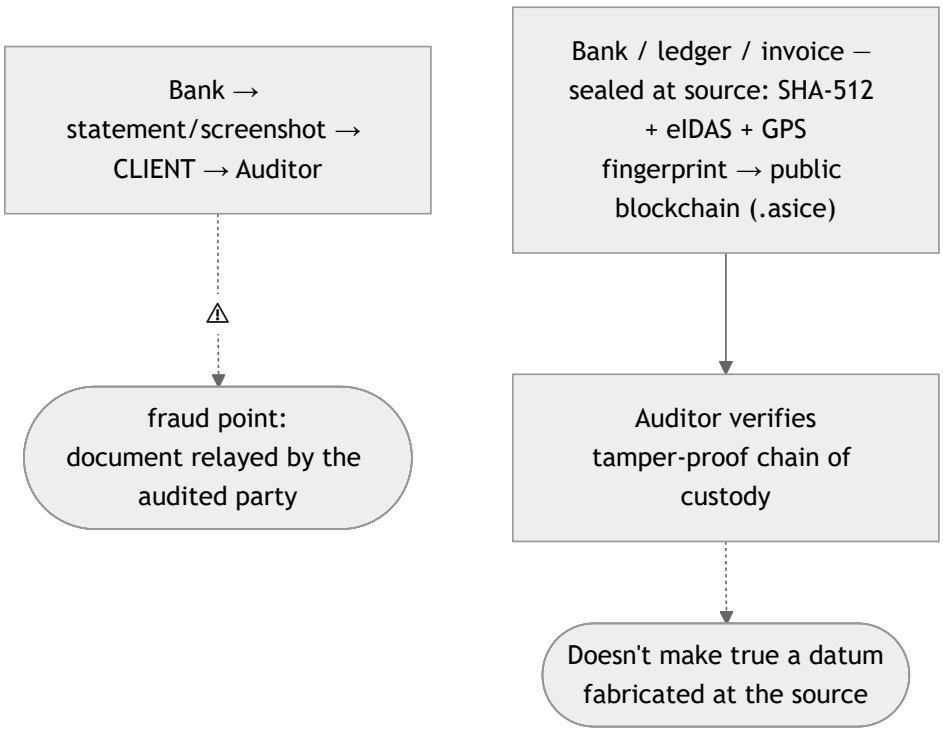


Figure 55.1 — In the vulnerable traditional model, the bank statement passes through the client (Wirecard) before reaching the auditor: that is the point where the forgery slips in. With t2p, the bank confirmation, the general ledger or the invoice is sealed at the source — SHA-512 fingerprint, eIDAS timestamp, GPS, anchoring to the public blockchain — and the auditor verifies an unalterable chain of custody (.asice). Limit: the seal does not make true a datum fabricated at the source (Wirecard, Luckin, Enron). Context: Wirecard, a nonexistent €1.9 billion; financial statement fraud, a median of $766,000 per case. Sources: Financial Times; PCAOB; ACFE.

---

[1] Financial Times (via S&P Global) — *Wirecard auditor did not independently verify its bank balance* (EY; nonexistent escrow accounts; fabricated statements). https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/wirecard-auditor-did-not-independently-verify-its-bank-balance-8211-financial-times-59219155 ↩

[2] CNN Business — *Germany bans EY from taking on new audit clients over Wirecard* (APAS: "grave negligence"; €500,000 fine and 2-year ban). https://www.cnn.com/2023/04/03/business/wirecard-ey-ban-germany/index.html

[3] Dan McCrum (Financial Times) — *Money Men: A Hot Startup, a Billion-Dollar Fraud, a Fight for the Truth* (2022). https://danmccrum.com/ ↩

[4] Wikipedia — *Enron scandal* (>$74 billion lost; Arthur Andersen, ~85,000 employees; origin of Sarbanes-Oxley). https://en.wikipedia.org/wiki/Enron_scandal ↩

[5] U.S. Securities and Exchange Commission — *Luckin Coffee agrees to pay $180 million penalty to settle accounting fraud charges* (~$311 million in fabricated sales). https://www.sec.gov/newsroom/press-releases/2020-319 ↩

[6] PCAOB — *AS 2310, The Auditor's Use of Confirmation* (adopted September 2023; direct external confirmation). https://pcaobus.org/oversight/standards/auditing-standards/details/AS2310 ↩

[7] Association of Certified Fraud Examiners (ACFE) — *Report to the Nations 2024* (5% of revenue; financial statement fraud, a median of $766,000 per case). https://www.acfe.com/-/media/files/acfe/pdfs/rttn/2024/2024-report-to-the-nations.pdf ↩

[8] European Commission — *Corporate sustainability reporting (CSRD)* (mandatory limited assurance; EFRAG's ESRS standards). https://finance.ec.europa.eu/financial-markets/company-reporting-and-auditing/company-reporting/corporate-sustainability-reporting_en ↩

[9] Accountancy Europe — *Omnibus explained: key changes to the CSRD and CSDDD* (threshold of >1,000 employees and >€450 million; limited assurance remains). https://accountancyeurope.eu/publications/omnibus-explained-key-changes-to-the-csrd-and-csddd/ ↩

[10] IAASB — *ISSA 5000: General Requirements for Sustainability Assurance Engagements* (2024; first global sustainability assurance standard). https://www.iaasb.org/publications/international-standard-sustainability-assurance-5000-general-requirements-sustainability-assurance ↩

Advertising Fraud: The Ad No One Saw

"Promise, large promise, is the soul of an advertisement."

— Samuel Johnson, The Idler, No. 40 (1759)

In October 2016, a bot farm run from Russia pretended to watch up to **three hundred million video ads a day**. No one was watching: the software simulated mouse cursors that moved, sessions logged into social networks and page visits; that is how it collected from the advertisers. The operation, dubbed Methbot, took in between three and five million dollars a day [6]. Its successor, a network called 3ve, went further: it infected **1.7 million home computers** without their owners knowing, impersonated ten thousand legitimate domains and cost advertisers more than **$29 million** in ads no human being ever saw [5]. Behind it was a Russian who called himself "the King of Fraud"; sentenced to ten years in prison, he wrote in his defense: "I am a soldier without weapons facing a tank called the FBI" [7].

That case illustrates a massive and surprisingly invisible problem: a good part of what we spend on digital advertising is lost on ads nobody sees. This chapter is about how to make the most basic question of the whole business provable again: **was an ad really served to a real person?**

The ad no one saw

The magnitude is chilling. The analysis firm Juniper Research calculates that digital advertising fraud cost $84 billion in 2023 — 22% of all online advertising spend, and up to 30% on mobile — and projects it will reach $172 billion in 2028 [1]. No surprise if we consider a hair-raising fact: according to the security firm Imperva, in 2023 bots accounted for 49.6% of internet traffic, and malicious bots, at 32%, grew for the fifth consecutive year and imitated human behavior to evade detection [4]. When half the traffic is not human, knowing who really sees an ad stops being a technical detail and becomes the sector's central problem.

The dollar that evaporates

But criminal fraud is only one part. The other is **the opacity of the chain**. The American advertisers' association, the **ANA**, led by Bob Liodice, published in December 2023 a study showing that, of every dollar an advertiser puts into the system, only **36 cents** actually reach the consumer; the rest vanishes in transaction costs and in "media productivity loss," a euphemism that includes so-called *made-for-advertising* sites — junk websites created solely to capture ads — which at one point took 15% of the spend [2]. The ANA calculates that advertisers could recover some **$22 billion** just by cleaning up their chain [2]. As the independent researcher **Augustine Fou**, one of the sector's most critical voices, warns, the platforms themselves give an incomplete picture of campaign success, and the fraud figures the industry reports tend to underestimate the problem.

...and the knot of it all is, once again, a problem of **proof**: no one can independently demonstrate that an impression was served to a real human, in a specific place and at a specific time. Transparency standards exist — the IAB Tech Lab's ads.txt, which declares which sellers are authorized — and viewability metrics, but they all rest on trust and on data controlled by whoever sells the inventory.

The proof of existence: the provable human eye

This is where **t2p** fills the gap. Every impression event generates a record — ad served, attested device identifier, viewability per the industry standard, geolocation — and the corresponding **SHA-512** is anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS. The record becomes unalterable: no one can inflate the figures afterward or rewrite how many times an ad was seen. Moreover, just as ads.txt declares *who* may sell, **t2p** can **chain and cryptographically seal every hop** of the chain — from publisher to sell-side platform, to exchange, to buy-side platform — so that every resale carries its verifiable fingerprint, time and origin. Thus, as we have seen, these lost amounts become a traceable chain. And to validate an influencer's real reach, a campaign with **proof of life** and **forensic GPS** can demonstrate that behind a view there is a living person in a real place, not a farm of fake followers — a fraud that costs brands hundreds of millions. Everything is packaged in an **.asice**

verifiable offline that an independent third party can audit without depending on whoever sells the inventory — precisely the "incomplete picture" Fou denounces.

The limit must be stated bluntly, because here it is threefold. The proof of existence demonstrates the integrity, time, place and — with proof of life — human presence of a record, but by itself it does not resolve three questions. The first, **privacy**: verifying that behind every impression there is a real human implies, by definition, collecting signals from the device or the person — the same tension between surveillance and transparency that runs through this whole book — and it demands minimizing data and attesting without identifying the user. The second, that **bots evolve**: Methbot was already faking cursors, and today's attestation will be evaded tomorrow; the proof reduces fraud, it does not eliminate it. And the third, that it **does not fix the misalignment of incentives**: as long as intermediaries are paid per impression served and not per value delivered, there will be pressure to inflate. "Garbage in, garbage on-chain": cryptography makes the record inviolable, but it does not guarantee that what gets recorded is honest.

The promise and the peril

The promise can be quantified in the tens of billions; a layer of proof of existence over impressions and the programmatic chain would make it harder to charge for invisible ads, would offer advertisers evidence auditable by a third party and would recover part of these figures, as we have seen. In a market where half the traffic is not human, proving the real human eye restores the foundation of the whole business.

The peril calls for the usual nuance. First, privacy: the fight against fraud must not become surveillance of the user. Second, overconfidence: no attestation is definitive against bots that innovate without pause. And third, the incentive: the technology records inviolably, but the honesty of the chain depends on those who want it. The proof of existence is a powerful tool for transparency, not a silver bullet against fraud or a substitute for fixing the sector's incentives.

It's time to act

For advertisers: demand sealed, auditable evidence for every impression, instead of trusting the inventory seller's figures. For the programmatic chain (sell- and buy-side platforms, exchanges): seal every hop so that every resale is traceable. For the standards bodies (IAB, viewability metrics): integrate the proof of existence as a reinforcement of ads.txt and viewability. And for the platforms: stop offering an incomplete picture and open up to independent verification.

A bot farm got paid for three hundred million videos no one watched. Giving every impression and every hop of the chain a sealed proof of the real human eye — while never forgetting privacy or incentives — is, on this front of the commerce of attention, time2act.

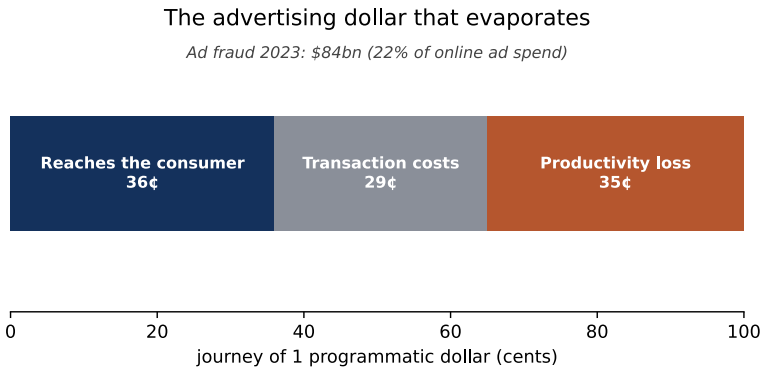


Figure 56.1 — Of every dollar an advertiser puts into programmatic buying, only 36 cents reach the consumer: 29 are lost in transaction costs and 35 in "media productivity loss" (junk websites and low-quality inventory). Digital ad fraud cost $84 billion in 2023 — 22% of online spend — and bots are already half of internet traffic. Sources: ANA (2023); Juniper Research; Imperva.

[1] Juniper Research (via PR Newswire) — *22% of online ad spend is wasted due to ad fraud in 2023* ($84 billion in 2023 → $172 billion in 2028; mobile 30%). https://www.prnewswire.com/news-releases/new-ad-fraud-study-22-of-online-ad-spend-is-

wasted-due-to-ad-fraud-in-2023-according-to-juniper-research-301938050.html ↩

[2] Association of National Advertisers (ANA) — *Programmatic Media Supply Chain Transparency Study* (December 2023; 36 cents per dollar; ~$22 billion recoverable; MFA 15%). https://www.ana.net/miccontent/show/id/rr-2023-12-ana-programmatic-media-supply-chain-transparency-study ↩

[3] eMarketer — *5 key stats on ad waste, fraud, MFAs and brand risks* (MFA from 15% to 4%). https://www.emarketer.com/content/5-key-stats-ad-waste-fraud-mfas-brand-risks-beyond

[4] Imperva — *2024 Bad Bot Report* (bots, 49.6% of internet traffic in 2023; malicious bots, 32%). https://www.imperva.com/company/press_releases/bots-make-up-half-of-all-internet-traffic-globally/ ↩

[5] Department of Justice (EDNY) — *Two international cybercriminal rings dismantled (Methbot and 3ve)* (November 2018; 1.7 million computers; >$29 million). https://www.justice.gov/usao-edny/pr/two-international-cybercriminal-rings-dismantled-and-eight-defendants-indicted-causing ↩

[6] Krebs on Security / White Ops — *Report: $3-5M in ad fraud daily from Methbot* (2016; 300 million views/day). https://krebsonsecurity.com/2016/12/report-3-5m-in-ad-fraud-daily-from-methbot/ ↩

[7] The Record (Recorded Future) — *"King of Fraud" sentenced to 10 years for Methbot/3ve* (Aleksandr Zhukov, 10-Nov-2021). https://therecord.media/king-of-fraud-sentenced-to-10-years-in-prison-for-role-in-methbot-3ve-botnet ↩

[8] IAB Tech Lab — *ads.txt (Authorized Digital Sellers)* (2017; against domain spoofing). https://iabtechlab.com/ads-txt/

[9] CBS News / Cheq (Roberto Cavazos, Univ. of Baltimore) — *Influencer marketing fraud costs companies $1.3 billion* (2019). https://www.cbsnews.com/news/influencer-marketing-fraud-costs-companies-1-3-billion/

[10] Augustine Fou (FouAnalytics) — independent research on digital advertising fraud. https://www.linkedin.com/pulse/how-do-we-un-fu-digital-advertising-dr-augustine-fou

Cybersecurity: Records the Attacker Cannot Rewrite

"Who controls the past controls the future: who controls the present controls the past."

— George Orwell, 1984 (1949)

When an attacker manages to get into a network, the first thing they do is not steal; they **erase their trail**. They run a command as simple as `wevtutil cl Security` to empty the Windows security log and remove the evidence of their passage. The maneuver has a catalogued name — technique **T1070, "Indicator Removal,"** from the MITRE ATT&CK framework — and it is so expected that defenders consider it a predictable step, not a surprise [2]. But it hides a revealing irony: on emptying that log, Windows automatically writes a new event, **Event ID 1102** — "the audit log was cleared" — with the account that did it [2]. Erasing the evidence *creates* evidence... **provided that evidence left the compromised system in time**.

And there lies the problem.

According to the annual report of Google and Mandiant's incident response unit, the average time an attacker operates inside a network without being detected — the *dwell time* — worsened to **fourteen days in 2025** [1]. Fourteen days are an eternity: time to spare for manipulating the records if they live only on the systems the attacker already controls. This chapter explains how to turn the history of an attack into a record that **not even the attacker can rewrite**.

The attacker attacks the evidence first

Cybersecurity has learned a hard lesson: the adversary attacks the proof before it attacks you. The consequences of not being able to prove what happened are ever more serious, because today it is not only the technicians who answer for it, but also regulators, courts and insurers. As we have seen, that year the global average cost of a data breach was $4.44 million, with $10.22 million in the United States and an average of 241 days to identify and contain it [4]. Regulatory pressure has soared: the European data protection regulation requires notifying a breach within 72 hours; the NIS2 directive and the DORA regulation impose strict incident notification deadlines; and the American securities commission requires disclosing a material incident within four business days [5][6][7][8].

The case of the city of Hamilton, Ontario, shows how far it goes. After a ransomware attack in February 2024 that paralyzed much of its network, the insurer denied coverage because the city had not fully deployed two-factor authentication, leaving taxpayers with a bill of up to 18.3 million Canadian dollars [9]. What the city's records and attestations could prove — or not — decided who paid. In cybersecurity, the integrity of the evidence is no longer just a technical question: it is a question of millions.

Getting the evidence out of the attacked system

The standard defense against log wiping consists of getting the logs off the system: **forwarding them in real time** to an external store or to a security event management system, so they survive the local purge. It is good practice, but it has a hole: a sufficiently skilled attacker can **disable the forwarding before acting**. And even if the records reach an external server, who guarantees that someone — an insider, or the attacker with privileged access — does not alter them there? A stronger guarantee is needed: that the recorded history is anchored in a place no one can rewrite. It is exactly what the reference forensic standard, the **NIST SP 800-86** guide, describes: the usefulness of digital evidence depends on being able to show it has not been altered, and on a chain of custody that holds up in court [3].

The proof of existence: records that cannot be rewritten

This is where **t2p** provides that anchor. A security operations center cannot anchor every log line — the volume would be unworkable; what is anchored periodically is the root of a Merkle tree summarizing a whole batch of records. **t2p** takes the **SHA-512 fingerprint** of that root, together with the batch header (time range, source, number of events), **seals it with an RFC 3161 timestamp** conforming to eIDAS and **anchors it to the public blockchain**. The chain of roots forms a history that can only be appended

to, never modified: each root incorporates the previous state, so rewriting a past event would require rewriting all the later roots already anchored on a public chain — impossible without the verification failing.

This frustrates exactly the technique described; if the attacker deletes or alters records *after* their fingerprint is already anchored, the discrepancy between the reconstructed log and the anchored root gives them away. And if they delete records that had **not yet** been anchored, the gap in the continuity of the chain of roots is evidence of tampering. For the evidence of a breach — captures, memory dumps, incident response chronologies — **t2p** generates an **.asice** container verifiable offline with fingerprint, timestamp, **forensic GPS** of the place of capture and the **proof of life** of the analyst who collected it: a chain of custody aligned with the standard that serves before a regulator, an insurer — the Hamilton case — and a court, proving *when*, *where* and *who* preserved each artifact, without depending on the victim's word.

The limit must be stated with full clarity, because here it is essential. The anchor **does not prevent the breach**: it does not stop the attacker or close the vulnerability. It only protects **what was recorded**; if the logging had holes, was misconfigured, or the attacker disabled the forwarding before acting, there is nothing to anchor — "garbage in, garbage on-chain." The integrity is only as good as the **immediacy of the forwarding** off the compromised system: if the root is anchored every hour and the attacker wipes inside that window before the sync, that stretch is unprotected. And there remains the performance cost: an operations center generates millions of events, and the anchoring frequency has to be chosen. **t2p** guarantees that **no one rewrote the record without leaving a trace**, not that the record is complete or that it reflects the whole truth of the attack.

The promise and the peril

So the promise is tangible: a layer of anchoring over the logs would turn the classic "your word against mine" into a chronology no one can touch, would render trail-wiping useless and would give the response team evidence admissible before regulators and insurers. It could also save the millions Hamilton lost; when a breach costs on average more than four million dollars, an unalterable proof of what happened becomes a first-order asset.

The peril calls for the usual nuance. First, do not confuse anchoring with prevention: sealing the logs stops no attack. Second, avoid overconfidence; the proof is only worth as much as the immediate forwarding and the quality of the logging. Third, protect the privacy of the logs themselves, which can contain sensitive data and need a design that anchors the fingerprint, not the cleartext content. The proof of existence is a powerful tool for the post-incident phase, but it does not replace defense or good security hygiene.

Furthermore, it must be acknowledged that the chain's immutability protects the integrity of the records without guaranteeing confidentiality. Once a file is uploaded to a public blockchain, any node can validate it and, given additional information, deanonymize it. American federal agencies such as ICE, the IRS or the FBI already buy chain analytics to link IP addresses to real identities, turning what should be anonymous evidence into a direct trail to specific people [12]. This scenario becomes especially dangerous when the payload is sensitive: medical documents, financial reports or personal data. An attacker who compromises the initial entry point can expose all the content, and immutability only ensures that the exploitation persists indefinitely. As Carissa Véliz warns, "it is possible that cybersecurity will never be achieved," because defenses evolve at the same speed as analysis and deanonymization techniques [11].

It's time to act

For security operations centers and their leaders: apply anchoring to the Merkle roots of the logs in near real time, so that trail-wiping stops working. For incident response teams: seal the evidence with a chain of custody and make it admissible before regulators, insurers and courts. For insurers: recognize anchored evidence as a basis for coverage decisions. And for regulators (NIS2, DORA, the securities commission): weigh the seal of origin as a reinforcement of the notification obligations.

An attacker tries to erase the logs to remove their trail, but the purge leaves a fingerprint — if it surfaces in time. Giving every record and every piece of evidence an unalterable anchor that turns the history of an attack into something no one can rewrite is, on this ever more expensive front, time2act.

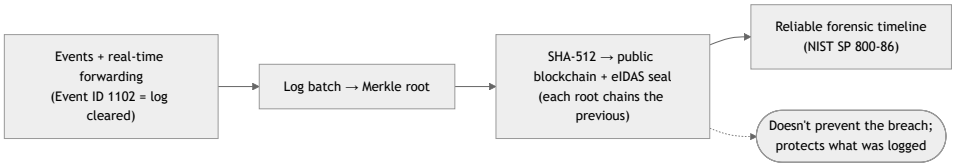


Figure 57.1 — System events (including Event ID 1102, "the log was cleared") are forwarded in real time off the attacked system; the log batches are summarized in a Merkle root that is sealed with a SHA-512 fingerprint and an RFC 3161/eIDAS timestamp and anchored to the public blockchain, chaining each root to the previous one (append-only). Any later alteration breaks the fingerprint and exposes the attacker. Limit: the anchor does not prevent the breach; it only protects what was recorded and forwarded in time. Context: average dwell time of 14 days in 2025; average cost of a breach, $4.44 million. Sources: MITRE ATT&CK; Mandiant M-Trends; NIST SP 800-86.

---

[1] Mandiant / Google Cloud — *M-Trends 2026* (average dwell time of 14 days in 2025; access handoff in 22 seconds). https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026 ↩

[2] MITRE ATT&CK — *T1070 Indicator Removal / T1070.001 Clear Windows Event Logs* (the purge leaves Event ID 1102). https://attack.mitre.org/techniques/T1070/001/ ↩

[3] NIST — *SP 800-86: Guide to Integrating Forensic Techniques into Incident Response* (2006; forensic phases and chain of custody). https://csrc.nist.gov/pubs/sp/800/86/final ↩

[4] IBM — *Cost of a Data Breach 2025* (global average cost of $4.44 million / $10.22 million in the U.S.; 241 days to identify and contain). https://www.ibm.com/think/x-force/2025-cost-of-a-data-breach-navigating-ai ↩

[5] European Commission — *NIS2 Directive (EU 2022/2555)* (notification within 24 h / 72 h; 23 states in infringement proceedings). https://digital-strategy.ec.europa.eu/en/policies/nis2-directive ↩

[6] EUR-Lex — *DORA, Regulation (EU) 2022/2554* (applicable since 17-Jan-2025; reporting of serious ICT incidents). https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng ↩

[7] U.S. Securities and Exchange Commission — *Cybersecurity disclosure rule* (26-Jul-2023; Form 8-K, item 1.05, within 4 business days). https://www.sec.gov/newsroom/press-releases/2023-139 ↩

[8] GDPR-info — *Article 33 GDPR* (data breach notification within 72 hours). https://gdpr-info.eu/art-33-gdpr/ ↩

[9] Global News — *Hamilton cyberattack cost* (insurer denies coverage over incomplete MFA; 18.3M CAD). https://globalnews.ca/news/11313018/hamilton-cyberattack-cost/ ↩

[10] Richard Bejtlich (TaoSecurity) — *Mandiant global median dwell time* (two consecutive years of deterioration). https://taosecurity.blogspot.com/2026/03/mandiant-global-median-dwell-time.html

[11] Carissa Véliz (ed.) — *The Oxford Handbook of Digital Ethics* (Oxford UP, 2022). ↩

[12] FedScoop (2024-2025). https://fedscoop.com/ice-wants-more-blockchain-analytics-tech/ ↩

Your Leaked Data: The Citizen's Claim

"Saying that you don't care about privacy because you have nothing to hide is no different from saying you don't care about freedom of speech because you have nothing to say."

— Edward Snowden, Permanent Record (2019)

The email opens with the usual line: "The security of your data is our priority." Three paragraphs down, in slightly smaller print, it admits that your name, your phone number and your ID number have been "potentially affected by an incident." And from that day on, the steady drizzle: the text message from "your bank" that is not your bank, the call that knows things it should not know, the attempt to get into your accounts, the SIM duplicate request you never made. The breach is not a one-off event: it is a drip that lasts months and years — and that almost no one can bill to whoever caused it.

The framework is there, halfway [2]: the GDPR requires companies to notify breaches to the authority within 72 hours and, if the risk to those affected is high, to inform them directly — the obligations chapter 57 explains from the company's side. And public tools such as Have I Been Pwned let anyone check whether their email appears in known leaks. What was missing was the economic piece: who pays for the harm this does?

This chapter is the citizen's mirror of chapter 57: there, the records of the attacked company; here, the life of the person whose data now circulates on the black market. Digital identity is in chapter 47; the scam that uses leaked data against the elderly, in chapter 85; SIM swapping, in chapter 68. The key question: can you collect for the harm of a leak? The answer: the European Court of Justice has said yes — with no minimum threshold of seriousness — but with three conditions that must be PROVEN [1]. And proving them is, exactly, a matter of chronology: the one built with **t2p** — fingerprint computed on the device and a certain time anchored to the public chain (chapter 13) [3] — from the day the notice arrives.

The CJEU opens the door — and sets an evidentiary threshold

On May 4, 2023, the Court of Justice of the European Union handed down judgment in case C-300/21 — the case of the Austrian postal service that, using an algorithm, inferred citizens' political affinity without their consent — and settled the doctrine of Article 82 of the GDPR [1]: the right to compensation requires three cumulative conditions — an infringement of the regulation, harm (material or non-material) and a causal link between the two.

The judgment contains two tensions that define this chapter. The one favorable to the citizen: non-material harm is compensable with no minimum threshold of seriousness — feeling exposed, humiliated or anxious counts. The hard one: the mere infringement is not enough — the harm must be proven, and so must the causation; and quantification is left to the national courts [1].

Europe has said, then, that your distress is worth money — if you prove it. And this is where the citizen always loses: the breach happens one day, the phishing arrives over months, the impersonation attempt a year later — and by the time the claim is filed, everything is scattered memories in deleted junk-mail folders. The practical reading is this book's: the causation the CJEU demands is, exactly, a chronology — and chronologies are built with certain dates, not with memory.

The proof of existence applied

The diary starts the day the notice arrives [3].

1. The notification captured

The email or the letter informing you of the breach — or the page where the company admits it — captured with a certain date [3]. It is the first link of the causation: when it became known, what was admitted, which data were acknowledged as affected. And if the company did not notify you when it should have [2], the absence is documented too: the date you learned of it from the press or from a public tool, sealed.

2. The certificate of presence

The check on Have I Been Pwned — your email, which leaks it appears in — captured with t2p's certified browser: the full image and the source code, with a date [3]. Your presence in the leak stops being your assertion: it is a dated query of a recognized public source.

3. The diary of the harm

Every consequence, sealed in the moment: the phishing text that quotes your full name, the call that knew your ID number — a recording of your own conversation, as chapter 75 explains —, the access-attempt alert, the SIM duplicate request you never made (chapter 68) [3]. Thirty seconds per piece. A year later, you do not have memories: you have a dated series that starts right after the breach — the causal link the CJEU demands [1], made of facts.

4. The claim that arrives armed

The whole file — notification, presence, diary of the harm — exported as an .asice package verifiable offline [3]: for the claim against the company, the complaint before the data protection authority or the civil suit under Article 82 [1][2]. And the usual symmetry: the company that notified properly and on time — and that has its chapter 57 records anchored — can prove its diligence in the same currency.

The limits, as always, upfront: the seal attests to the chronology — not the valuation of the harm, which belongs to the judge with national law in hand [1]. Correlation is not always causation: phishing existed before the breach, and the expert will have to assess the connection — the dated series helps them, it does not replace them. And the diary of the harm is yours: on the device, only fingerprints on the chain, selective disclosure [3].

The promise and the peril

The promise: citizens who claim with chronologies instead of complaints [1][3]; companies that respond with provable diligence instead of templates; the non-material harm of leaks — recognized by the CJEU with no threshold [1] — turned into viable claims; and a new, powerful incentive to protect data beforehand: when the harm can be proven, negligence gets expensive.

The perils, which are not few.

The first is **the claims industry**: firms mass-producing identical suits with template harms. The counterweight is the doctrine itself: the CJEU requires proven harm and causation [1] — the template claim dies where the personal chronology begins.

The second is **documentary revictimization**. Asking the victim to seal everything can be a second burden. The gesture must be minimal — the ten-second capture — and the tool, an ally: those who do not document lose none of their rights, only their best evidence.

The third is **the diary as new loot**: the file on your harm contains... more of your data. Hence the golden rule: on the device, only fingerprints on the chain [3], and selective disclosure to whoever is entitled.

The fourth peril is confusing redress with protection. No seal prevents the breach or puts the data back in the jar. Prevention is the task of companies — chapter 57 — and of the regulator; the proof of existence only ensures that, when they fail, the cost is not paid in silence by the victim.

It's time to act

The diary starts on day one — because day one thousand will be too late.

For you, citizen: when the notice of a breach arrives — or the suspicion — start the diary: the notification, the presence, every symptom dated [3]; tomorrow's claim is built on day one.

For consumer associations and law firms: demand and submit anchored chronologies — the CJEU has given you the door [1]; the key is dated causation, not volume.

For the companies that hold data: notify properly and anchor it — your provable diligence is the difference between an incident and a conviction; chapter 57 is yours.

For the data protection authorities: weigh the citizen's dated file as what it is — causation made chronology [1]; and publish guides so documenting the harm is within everyone's reach.

For legislators: Article 82 is already there [1][2]; the pending step is procedural — recognize the anchored chronology as qualified evidence of non-material harm, so that claiming does not require being an expert.

Your data may travel through places you will never see; let at least the harm they do to you leave a trail that you will. For the privacy that is defended with facts, **time2act**.

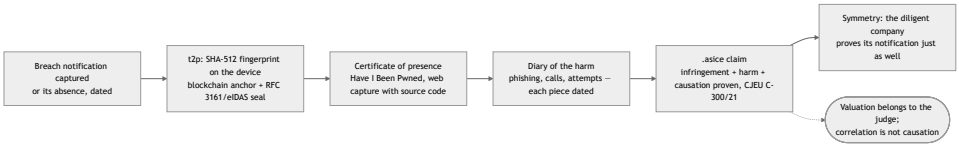


Figure 58.1 — The breach seen from the citizen's side, with proof of existence (sealed with t2p — chapter 13): the notification captured (or its absence, dated), the certificate of presence in the leak (web capture with source code), the diary of the harm (phishing, calls, attempts — each piece sealed in the moment) and the .asice claim. Foundation: CJEU C-300/21 — non-material harm compensable with no threshold, but with infringement + harm + causation PROVEN. Symmetry: the diligent company proves its notification. Limits: valuation belongs to the judge; correlation ≠ causation; the diary stays on the device. Sources: CJEU; GDPR Art. 82.

---

[1] CJEU — *Judgment of May 4, 2023, case C-300/21 (UI v. Österreichische Post)* (Article 82 GDPR: three cumulative conditions — infringement, harm and causal link; the mere infringement confers no right to compensation; non-material harm is compensable with no minimum threshold of seriousness; quantification under national law subject to equivalence and effectiveness); analyses by Cuatrecasas, finReg360 and Confilegal. https://www.cuatrecasas.com/es/spain/art/indemnizacion-de-danos-morales-por-infraccion-del-rgpd · https://finreg360.com/alerta/el-tjue-dicta-sentencia-sobre-el-derecho-de-indemnizacion-por-tratamiento-de-datos-personales/ · https://confilegal.com/20230913-el-tjue-admite-que-los-danos-no-materiales-por-el-uso-ilicito-de-datos-personales-merecen-una-compensacion-economica/ · https://www.iberley.es/noticias/tjue-mero-incumplimiento-normativa-proteccion-datos-da-lugar-responsabilidad-pero-solo-dano-probado-resulta-indemnizable-33163 ↩

[2] GDPR (Arts. 33-34: breach notification to the authority within 72 h and to those affected in case of high risk; Art. 82: right to compensation); Have I Been Pwned as a public leak-checking tool. See chapter 57 for the company's side. ↩

[3] t2p — product manual (certified screenshot and web page capture — full image + source code; audio recording of one's own conversations; certify any file; projects for an incident diary; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; SIM swapping, in chapter 68; the scam with leaked data, in chapter 85. ↩

The Documented Outage: SLAs, Cloud and Continuity

"No man is an island, entire of itself; every man is a piece of the continent, a part of the main."

— John Donne, Devotions, Meditation XVII (1624)

On July 19, 2024, a faulty update to CrowdStrike's Falcon Sensor security software knocked out some 8.5 million Windows computers worldwide — the biggest IT outage in history [1]. Blue screens on airport terminals, in operating rooms, in bank branches, in railway systems. Delta Air Lines canceled more than 7,000 flights over five days, with more than 1.3 million passengers affected — and in October it sued CrowdStrike claiming some $550 million; the two companies ended up suing each other [1].

That day, thousands of companies discovered at once three things they had taken for granted. That their whole business hangs on one provider. That the contract promises "99.9% availability" with service credits as the only remedy. And that, when it is time to claim, the outage must be proven by the customer — on the status pages, dashboards and measurements of the very provider that failed. Who went down, when exactly, for how long, with what impact? The official answers are written by one of the parties.

This chapter deals with availability — security's forgotten sibling. The records an attacker cannot rewrite are in chapter 57; telecommunications and voice, in chapter 68; the insurance claim, in chapter 28. The key question: whose clock is it when the cloud goes down? The answer: the one who caused it — unless the customer has their own, with an opposable certain date: independent monitoring sealed with **t2p** — fingerprint, certain time and public anchor, the machinery of chapter 13 [3]. And with the usual symmetry: the provider that delivers has, in its own anchored records, the best commercial argument in the sector.

The provider's clock

The anatomy of the problem, piece by piece. The SLA — the service level agreement — promises a percentage of availability and offers credits when it is breached; but the credits are not automatic: they require the customer to claim within a deadline, and the burden of proof is theirs. The official measurement belongs to the provider: its status pages — which during incidents often lag behind reality —, its dashboards, its metrics. And the customer's measurement tools, to make matters worse, often run in the same cloud that is sinking.

The credit, moreover, covers a fraction of the monthly bill. The real damage — the lost sales, the stalled point of sale, the reputation — falls outside, and requires expert proof of lost profit. Without a reliable, verifiable chronology, the negotiation with the provider and with the insurer is a game of versions.

And the regulatory context raises the bar: DORA requires the European financial sector to manage the ICT risk of its providers and to report incidents precisely, and NIS2 extends the notification obligations to many more sectors — both ask, at bottom, for reliable chronologies [2]. Timekeeping is no longer just a contractual matter: it is a legal requirement.

The conclusion is this book's: the outage dispute is a dispute of clocks — and today almost all the admissible clocks belong to the same side. What is missing is not more monitoring: it is that the customer's monitoring have an opposable certain date.

The proof of existence applied

The moment the service goes down is the one instant the customer cannot afford to lose [3].

1. The independent probe, sealed

The company sets up its probes — periodic checks of the service from outside the monitored cloud — and seals their results with t2p: the log file of each interval, certified by fingerprint with a certain date; only the fingerprint is anchored, the log stays home [3]. The cadence is each business's choice; what matters is the chain: when the outage comes, the sequence "responding — responding — not responding — not responding" carries timestamps neither party can move. The customer's clock, at last, exists.

2. The status page, frozen

During the incident, the provider's status page is captured with t2p's certified browser — the full image and the source code, packaged with a certain date [3]. What the provider was saying at 10:42 — "all operational" or "incident under investigation" — is frozen. If the official narrative later changes, the discrepancy is a dated fact, not a memory.

3. The impact documented

The failing transaction, screen-recorded and certified [3]: the cart that will not charge, the unresponsive point of sale, the stalled order queue. The business metric exports — sales per hour, completed transactions — sealed by fingerprint. When lost profit has to be quantified before the insurer or the court, the expert will work on an anchored chronology, not on reconstructions.

4. The claim that stands

The .asice package of the incident [3]: probes, status captures, impact, communications with the provider — each piece verifiable offline and without trusting anyone. For the SLA credits, for the cyber insurance, for the DORA or NIS2 regulatory report [2]: the same evidence serves every front. And the symmetry: the serious provider that publishes its own anchored availability records turns its SLA into its best argument — against the opportunistic claim, its seals; against the competition, its transparency.

The limits, upfront: the seal attests to when the service was not responding from where the probe measured — it does not diagnose the cause or apportion blame among provider, network and customer: that is experts' work. A badly designed probe seals its own confusion — methodology first, cryptography second. And no fingerprint brings a downed service back up: continuity is built with redundancy and contingency plans — the proof only ensures that the damage, when it comes, is not also a dispute of chronologies.

The promise and the peril

The promise: SLA credits claimed with one's own chronologies [3]; disputes like those of July 2024 [1] that could be settled on both parties' anchored records instead of years of litigation; insurers and regulators [2] working on dated facts; and providers competing by publishing their verifiable availability.

And now the perils, unsweetened.

The first is **the war of probes**: each side measuring from wherever suits it best. The answer is methodological — probes from representative points, methodology declared and sealed in advance, like the pre-commitment of chapter 45, and experts who assess the quality of the measurement, not just its date.

The second is **the provider on the defensive**. If the customer's trail serves only to litigate, the relationship sours. The way out is symmetry: the provider that anchors its records earns the presumption of seriousness — verifiable transparency as the sector's standard, not as a weapon.

The third is **regulatory dead letter**. DORA and NIS2 [2] demand reports; if the reports are fabricated after the fact, the whole edifice is decorative. The anchored chronology is the minimum for compliance to mean anything.

The fourth is brutal in its simplicity: the proof does not bring servers back up. No seal replaces redundancy, multi-provider setups or the contingency plan. First, let the business survive the outage; then, let it be able to prove it.

It's time to act

Let the next outage find every company with a clock of its own.

For the companies that live in the cloud: set up the independent probe and seal its chain [3]; on the day of the next big outage [1], the chronology will be yours.

For cloud and hosting providers: anchor your availability records and publish them by fingerprint — the verifiable SLA is marketing no competitor can copy without delivering.

For cyber and business-interruption insurers: require anchored chronologies in the policy; you will pay real claims faster and fabricated ones less.

For financial and cybersecurity regulators: DORA and NIS2 [2] already demand incident chronologies — let them be verifiable, not narrative.

For legislators: SLA credits with unilateral measurement are a lame clause; recognizing independent measurement with a certain date would rebalance the most asymmetric contract in the digital economy.

The next global blue screen will come. Let it find every company with its sealed clock — and let the argument be about causes, not chronologies. For the economy that cannot stop, **time2act**.

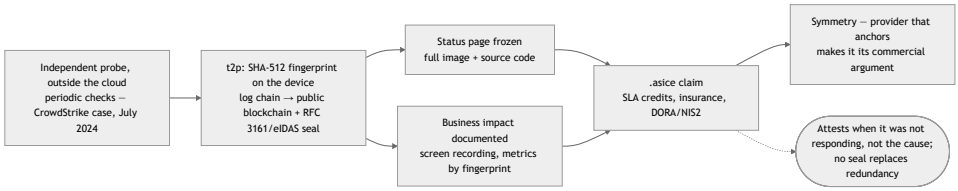


Figure 59.1 — The documented outage with proof of existence (sealed with t2p — chapter 13): the independent probe with its chain of sealed logs (the customer's clock), the provider's status page frozen with source code, the business impact documented (screen recording, metrics by fingerprint) and the .asice claim — for the SLA credits, the insurer and the DORA/NIS2 reports. Symmetry: the provider that anchors its availability makes it its commercial argument. Limits: the probe attests to when it was not responding, not the cause; methodology first; no seal replaces redundancy. Sources: CrowdStrike 19-7-2024 (8.5 million Windows machines; Delta ~$550 million).

---

[1] Wikipedia — *2024 Delta Air Lines disruption* / TechTarget — *CrowdStrike outage explained* (19-7-2024: faulty Falcon Sensor update; ~8.5 million Windows machines; the biggest IT outage in history; impact on airlines, hospitals, banks, emergency services); CNBC — *Delta, CrowdStrike sue each other over widespread IT outage* (Delta: >7,000 flights canceled over five days, >1.3 million passengers; October 2024 suit for ~$550 million; cross-suits). https://en.wikipedia.org/wiki/2024_Delta_Air_Lines_disruption · https://www.techtarget.com/whatis/feature/Explaining-the-largest-IT-outage-in-history-and-whats-next · https://www.cnbc.com/2024/10/25/delta-suit-against-crowdstrike-after-it-outage-caused-cancellations.html ↩

[2] Regulation (EU) 2022/2554 (DORA — digital operational resilience of the financial sector: third-party ICT risk management, incident reporting; applicable since 11-1-2025) and the NIS2 Directive (expanded incident notification obligations). See chapter 57 for the security context and notification deadlines. ↩

[3] t2p — product manual (certify any file — only the fingerprint travels; web page capture with full image + source code; certified screenshot and screen recording; projects to group an incident's file; background certification; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the records of an attack, in chapter 57; the insurance claim, in chapter 28. ↩

Art and Antiquities: Provenance Against Forgery and Plunder

"Even the most perfect reproduction of a work of art is lacking in one element: its presence in time and space, its unique existence at the place where it happens to be."

— Walter Benjamin, The Work of Art in the Age of Mechanical Reproduction (1935)

In a garage in the Queens borough of New York, an immigrant Chinese painter named Pei-Shen Qian manufactured masterpieces. He painted "Rothkos" and "Pollocks" that he aged with tea and vacuum-cleaner dust to simulate decades of grime, and he charged less than **nine thousand dollars** per canvas. New York's oldest gallery, Knoedler & Company — founded in 1846 — sold them for millions to collectors who trusted its one-hundred-and-sixty-five-year prestige. The fraud, worth some **eighty million dollars** across some forty forgeries, fooled the experts for almost fifteen years, until the gallery abruptly closed its doors in 2011 [1].

Now, how was the greatest German forger of the postwar era unmasked? Not by an expert or a document, but by a **can of paint**. The forger Wolfgang Beltracchi had sold a supposed Campendonk for €2.88 million, but laboratory analysis detected in it **titanium white**, a pigment that did not exist in the artist's lifetime [2]. That molecule set off a run of hundreds of forgeries. The lesson, which we put upfront because it is the heart of this chapter, is this: what authenticates a work is **science and the expert eye**, not paper. But what decides whether it can be sold, stolen or restituted is something else: its **provenance**, its history of origin and ownership. And there, the lie roams free.

The lie about the origin

If forgery attacks *who* made a work, plunder attacks *where it comes from*. The most emblematic case is that of the dealer Subhash Kapoor, owner of a Manhattan gallery, who ran an international antiquities trafficking network valued at some hundred million dollars. One of his objects was a bronze dancing Shiva from the eleventh century, torn from a disused temple in Tamil Nadu and sold in 2008 to the National Gallery of Australia for $5.6 million with false provenance papers. The Manhattan district attorney's antitrafficking unit ended up seizing more than 2,500 pieces linked to Kapoor, valued at some $143 million, and he was sentenced to ten years in prison in India in 2022 [3]. As the forensic archaeologist Christos Tsirogiannis, of the University of Cambridge, explains — he has identified thousands of plundered antiquities by cross-checking them against archives seized from traffickers — the "provenance" that auction houses and museums display is often a fabricated tale that begins right after the plunder: the problem is not the lack of papers, but that the papers lie.

The magnitude is as large as it is opaque. INTERPOL estimates the illicit trade in antiquities at between $225 million and $6 billion a year — the enormous range reflects the market's opacity [8]. In Syria, as we have seen, tens of thousands of pieces have been stolen from sites and museums. And the problem has a historical version that still stings: Nazi-looted art, like the 1,406 works — Picasso, Matisse, Chagall — that a routine customs check discovered in 2012 hidden in the Munich apartment of Cornelius Gurlitt, son of one of Hitler's dealers [5]. Against this, the international community has built defenses: the 1970 UNESCO Convention against illicit trafficking, ratified by close to 150 states [7]; the 1998 Washington Principles on the restitution of Nazi-looted art [6]; and databases such as the Art Loss Register, founded by Julian Radcliffe in 1991, which has some 700,000 stolen objects on record and has recovered more than a hundred million dollars' worth [9]. Radcliffe's thesis is luminous: if all stolen art is registered in a single place and everyone who buys or sells consults it, stealing stops being profitable because the piece cannot be resold cleanly.

The proof of existence: provenance from second zero

This is where **t2p** supplies what is missing: a **chain of provenance** as a succession of verifiable events. Each new link — a certificate of authenticity, a condition report, an invoice, an expert's opinion, a change of owner, a high-resolution photograph — is summarized in a **SHA-512 fingerprint** on the device, anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS and packaged in an **.asice container verifiable offline**.

But the most powerful application is archaeological. On a legal dig, the archaeologist can seal *in situ* the photograph of the **exact find-spot** — with forensic GPS and a timestamp — before moving the piece. This creates digital proof of a regulated excavation that distinguishes legitimate discovery from plunder: it is exactly the difference between Kapoor's Shiva, with papers invented years later, and an object with its **origin documented from second zero**. And the **proof of life** binds the expert or the excavator who signs the report to a real, living person at that moment, hindering the "phantom experts" and the certificates signed after the fact that appear in cases like Knoedler's. Against the pattern of Kapoor and Tsirogiannis, a sealed, dated record makes inventing a history *after* the moment of sealing detectable.

The limit must be said in all its rawness, because here it is absolute. The proof of existence **does not authenticate the object**. Sealing the photograph of a Beltracchi does not turn it into an authentic Campendonk: a sealed photo of a forgery is still a forgery. Authentication still demands the expert eye and scientific analysis — the pigments, the dating — which is what really caught Beltracchi. And, even more important, **any provenance fabricated *before* the seal stays fabricated forever**: the chain inherits and petrifies the lie. "Garbage in, garbage on-chain." The real value of t2p is not sealing old, dubious histories, but sealing the origin **at the exact moment** the evidence is born — the find-spot of an excavation, the report of an identified expert — so that from that instant on, the lie can no longer slip in.

The promise and the peril

The promise is great for a market eaten away by opacity. A chain of provenance sealed from the origin would make plunder much harder — because an antiquity without a documented find-spot would be immediately suspect — would hinder provenances "laundered" after the fact and would give museums, buyers and judges incontestable evidence of a work's history. In a trade where robbing a temple can yield millions, verifiable provenance is the best defense of humanity's heritage.

The peril calls for the usual nuance, and here it matters especially. The first, already said: confusing provenance with authenticity would be fatal; the seal makes nothing authentic, and science must remain irreplaceable. The second, that the chain petrifies whatever it is fed: a false provenance sealed today will be an immutable lie tomorrow. And the third, that the part that truly matters — the physical inspection, the expert's eye, the historical context — lives off-chain. The proof of existence is a powerful tool for documenting the origin from second zero, not a certificate of authenticity or a substitute for expert knowledge.

It's time to act

For museums and auction houses: demand sealed provenance and, above all, find-spots documented from the excavation, avoiding pieces whose histories begin suspiciously late. For archaeologists and legal digs: seal the find-spot in situ to generate proof that distinguishes discovery from plunder. For dealers and buyers: consult registries such as the Art Loss Register and do not trust deceptive documents. For regulators and UNESCO: recognize the sealed chain of provenance as the standard of legitimate trade. And for everyone, remember that no seal EVER replaces the expert's eye or the laboratory.

A can of titanium white exposed Europe's most prominent forger; false documents allowed a stolen god to be sold for millions. Giving every work a provenance sealed from its origin — never forgetting that authentication requires science, not just a fingerprint — is the front that guards memory and beauty, time2act.

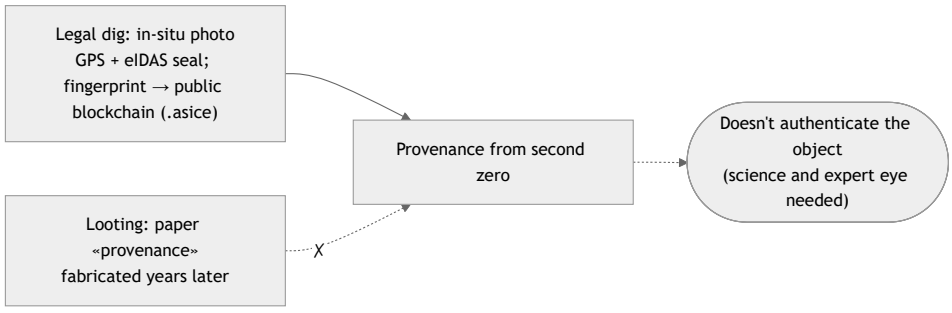


Figure 60.1 — On a legal dig (left), the piece is photographed in situ with GPS coordinates and an eIDAS timestamp before being moved: the provenance is documented from second zero and anchored to the public blockchain (.asice). In plunder (right), the paper "provenance" begins years after the find and is easy to fabricate — like the Chola Shiva sold to Australia for $5.6 million with false papers. Limit: the seal proves integrity, date, place and person, but does NOT authenticate the object (a sealed photo of a forgery is still a forgery). Context: illicit antiquities trade of $225 million to $6 billion/year. Sources: INTERPOL; Manhattan district attorney's office; Art Loss Register.

---

[1] CNBC — *This $80 million art scam fooled experts* (Knoedler gallery; Pei-Shen Qian; <$9,000 per work). https://www.cnbc.com/2018/08/16/think-you-can-spot-a-fraud-this-80-million-art-scam-fooled-experts.html ↩

[2] The Art Newspaper — *How the detection of titanium white gave away Beltracchi* (2012; the impossible pigment). https://www.theartnewspaper.com/2012/12/01/how-the-detection-of-titanium-white-gave-away-beltracchi ↩

[3] OCCRP — *India sentences notorious antiquities trafficker from Manhattan* (Subhash Kapoor; >2,500 pieces, ~$143 million; the Chola Shiva). https://www.occrp.org/en/news/india-sentences-notorious-antiquities-trafficker-from-manhattan ↩

[4] Manhattan District Attorney's Office — *Antiquities Trafficking Unit* (~6,000 antiquities from 29 countries, >$460 million recovered). https://manhattanda.org/d-a-bragg-announces-return-of-antiquities-to-spain-and-cyprus/

[5] The Conversation — *The Gurlitt case: how a routine customs check uncovered a Nazi-era art hoard* (1,406 works). https://theconversation.com/the-gurlitt-case-how-a-routine-customs-check-uncovered-a-sensational-nazi-era-art-hoard-105604 ↩

[6] U.S. Department of State — *Washington Conference Principles on Nazi-Confiscated Art* (1998; 44 countries). https://www.state.gov/washington-conference-principles-on-nazi-confiscated-art ↩

[7] UNESCO — *1970 Convention against the illicit trafficking of cultural property* (~149 states). https://www.unesco.org/en/fight-illicit-trafficking/about ↩

[8] INTERPOL — *Cultural heritage crime: the issues* (illicit trade of $225 million to $6 billion/year). https://www.interpol.int/en/Crimes/Cultural-heritage-crime/The-issues-cultural-property ↩

[9] Art Loss Register — *~700,000 registered objects; founded by Julian Radcliffe (1991); >$100 million recovered.* https://www.artloss.com/ ↩

[10] CNN — *The forensic archaeologist hunting trafficked antiquities* (Christos Tsirogiannis; "the papers lie"). https://edition.cnn.com/2025/12/13/style/forensic-archaeologist-trafficked-antiquities-scli-intl

Historical Memory: The Archives Against Denialism

"Those who cannot remember the past are condemned to repeat it."

— George Santayana, The Life of Reason (1905)

In June 2024, UNESCO and the World Jewish Congress published a report that should keep any historian up at night: generative AI threatens the memory of the Holocaust. It is not an abstract warning — it comes with examples: ChatGPT had invented campaigns of a "Holocaust by drowning," with Nazis drowning Jews in rivers, an episode that never existed; Bard had fabricated false quotes from testimonies to prop up distorted narratives of real massacres [1]. Consider for a moment what that means: the machine does not merely get things wrong — it invents atrocities with the appearance of history and testimonies with the appearance of a human voice, on an industrial scale and in the tone of an encyclopedia. And meanwhile, the last survivors are fading: soon, all living memory will be archive — and the archive, attackable.

Up to here this book has defended the present: proving what happens now, as it happens. This chapter extends the same logic backward. The provenance of art objects is dealt with in chapter 60; the photography of the present, in chapter 17; the general crisis of truth, in chapter 10. Here the question is different: how do you protect the past, when fabricating "evidence" of it costs one prompt? The answer: by anchoring memory before the flood arrives — archives with sealed manifests, testimonies fingerprinted while the witness lives, exhumations with an anchored chain of custody and photographic collections with a certain date, with **t2p**: the SHA-512 fingerprint, the RFC 3161/eIDAS timestamp and the anchor to the public blockchain, inside an .asice container [5]. The past cannot be changed; its proof can — and that is why it must be sealed.

History, rewritten by machine

The UNESCO report does not stop at the examples: it makes a structural diagnosis. Generative models tend to oversimplify historical complexity and prioritize a small selection of sources and events, leaving out whole nuances and contexts; the recommendations point in two directions — that the technology companies work with survivors, descendants, educators and historians, and that education systems equip young people with digital literacy and critical thinking [1].

And the phenomenon already speaks Catalan and Spanish. In February 2025, Maldita.es documented how websites and magazines were publishing on their front pages AI-generated images of real people and events — the pope, the scientist Margarita Salas, scenes from the Spanish Civil War — often without any label identifying them as synthetic; the risk described: altering the perception of historical figures and events, generating false narratives and deepening distrust in the media [2]. The fabricated "period" image is no longer a laboratory exercise: it is editorial content in circulation.

The mechanics of the damage is the liar's dividend applied to history. When fabricated period images are indistinguishable from authentic ones, the worst of both worlds happens at once: the false photo finds believers, and the true photo finds denialists — "that's a fake" becomes the universal reply. And denialism, which had always had to work against the documents, can now fabricate favorable documents. The battle for memory stops being merely interpretive and becomes evidentiary: which documents existed, since when, and who can prove it.

The archive that trembled

On October 9, 2024, it emerged that the Internet Archive — the digital library that guards the memory of the web — had suffered a major intrusion: 7 terabytes stolen, including the institution's source code and the addresses and hashed passwords of 31 million users, via a GitLab token that had been exposed since late 2022. Simultaneously, denial-of-service attacks knocked the Wayback Machine offline, and a second breach the same month exposed more than 800,000 support tickets, some with identity documents. The service came back in mid-October, with limited functions [3].

Why is this more than a technical incident? Because the Wayback Machine is the de facto memory of the internet: the tool with which journalists, judges and researchers prove that a page said one thing before someone quietly edited it. Every judicial or academic citation of a capture trusts the integrity and

availability of an infrastructure sustained by a single nonprofit — attackable, as was seen, through one forgotten token and one avalanche of traffic.

The lesson is not to blame the Archive, a venerable and essential institution: it is to understand that no institution of memory, however solid, can guarantee alone the integrity of what it guards. The structural answer is anchored manifests: the fingerprint of every collection and every capture, with a timestamp, published on a chain that no attack on any server can rewrite [5]. If the contents are lost, the existence and integrity of what was there remain attested; if someone alters them, the discrepancy sings. Memory needs copies; trust in memory needs anchors.

The graves that still speak

The Spanish front of this chapter is called Law 14/2022, on Democratic Memory: the central state administration took on the responsibility of searching for the people who disappeared during the Civil War and the dictatorship, the State DNA Bank of victims was created, and the Integrated Map of Graves was pushed forward, recording 5,848 located graves, of which around 1,500 have been opened [4]. The balance known in November 2025: 8,941 remains located and recovered and some 9,000 victims redressed — but only 70 remains fully identified and returned to their families, and more than 10,000 bodies still to recover [4]. Every figure is a family waiting.

This book's angle is the evidentiary one, and it is rigorously neutral. An exhumation is forensic science — stratigraphy, photographs, objects, samples, DNA — and its strength depends on the chain of custody. And it is, at the same time, an act exposed to three threats: future revisionist doubt ("these remains are not what they say"), the accusation of political manipulation, and human error.

Anchored documentation protects the process from all three. Every stage of the excavation can be sealed in the moment — the photograph of each find with a certain date, the sample log, the chain of custody to the State DNA Bank [4][5] — and the result is a file that no future denialism will be able to dismiss as "fabricated" without denying the mathematics, and that no suspicion of manipulation will be able to stain without evidence. The seal serves no memory against another: it serves the facts against oblivion and against the lie, wherever they come from.

The proof of existence applied

Any county archive, an afternoon of digitization: the sound of the scanners, yellowed pages of an old newspaper — and each sheet, as it is born digital, receives its certain date.

1. The anchored archive

Libraries, newspaper archives and archives — public, community, exile — publish anchored manifests of their collections: the SHA-512 fingerprint of each document and each batch, with a timestamp, on the public blockchain [5]. The silent editing of an old news item becomes detectable by anyone; the loss of a server does not erase the evidence of what existed [3]; and the academic or judicial citation of a document can be verified against the original fingerprint. It is the same logic the book applied to official statements (chapter 86), extended into deep time: not one single, official archive, but all archives, each with its integrity attested. When a freshly fabricated "article from 1940" appears, its absence from the anchored corpus will speak for it.

2. The sealed testimony

The generation of the survivors is ending, and UNESCO recommends working with them while they are here [1]. Every recorded testimony — voice, video, transcript — can be sealed within the witness's lifetime: a fingerprint and a certain date attesting that those words existed, exactly so, at that moment [5]. When in twenty years a denialist says a testimony "was fabricated after the fact," or when an AI generates false testimonies with a cloned voice [1], the difference will be verifiable: the real testimonies will have anchored temporal priority; the fabricated ones will not. And the same goes for the oral memory of the Civil War: every oral history interview sealed today is an armored document tomorrow.

3. The documented grave

Every exhumation under Law 14/2022 [4] can anchor its file: photographs of each stratum with a certain date, the forensic team's reports, the find log, the chain of custody of the samples to the State DNA Bank [4][5]. The result serves three causes at once: the judicial one (an impeccable chain of custody), the family one (a traceable identification — so that the 70 families who have already received their remains [4] become thousands, with the same certainty) and the historical one (a record that says, forever and with mathematics behind it: here there was a grave, this is what was found in it, and this is its record).

4. The pre-AI horizon

There is a window that is closing: everything digitized and anchored before the maturity of synthetic imagery is attested as prior — a pre-AI horizon separating the verifiable photographic corpus from everything that comes after. Museums, photographic archives and family collections that anchor their digitizations today [5] create the reference canon against which future fabricated "period" images [2] can be checked: if a photo "from 1937" appears nowhere in the anchored corpus and surfaces a few years from now, its lack of history will speak for it. Every year that passes without anchoring, the horizon blurs a little more.

The limits, stated clearly: the seal attests to when a document was digitized or recorded and that it has not been altered since — it does not attest that the document tells the truth, nor does it replace source criticism, which remains the historians' craft; it cannot recover what was never documented; and history is also interpretation — the seal protects the documents and the facts, not any single narrative about the facts.

The promise and the peril

The promise is a collective memory with verifiable integrity: archives that cannot be silently edited [3], testimonies with temporal priority over their forgeries [1], graves with unassailable files [4] and a pre-AI photographic canon against fabricated history [2]. The truth of the past would stop depending on a single voice or a single server: it would be validated against a fabric of fingerprints open to anyone.

The perils demand the same clarity.

The first is **the single archive**. If anchoring becomes the monopoly of an "official archive of truth," we will have created the ministry Orwell feared. The answer is plurality: all memories — institutional, community, exile, dissident — must be able to anchor their holdings. The seal attests to integrity and date; never authority, never a single version.

The second is **memory against personal oblivion**. Archives contain data on living people and families. Anchoring is not publishing: only fingerprints on the chain, with access governed by each archive's frameworks and by the rights of the people mentioned. Immutability protects the integrity of the document; it cannot be used to deny anyone their rights.

The third is **the fetish of immutability**. History corrects itself: new documents appear and old ones are reinterpreted. Anchoring documents versions and dates; it does not freeze knowledge. Confusing "sealed" with "definitive" would betray the very craft it is meant to protect — each new finding is anchored as a new version, and the debate goes on.

And the fourth is **memorial techno-solutionism**. No blockchain teaches history. Against denialism, the first line is school, historians and the critical literacy UNESCO calls for [1]; the seal only guarantees that, when it is necessary to prove what existed, it can be done. A society that does not want to remember will not be saved by cryptography.

It's time to act

The past cannot be changed; its proof can — and that is why it must be sealed now.

For archives, libraries and newspaper archives: publish anchored manifests of your collections [3] [5]; let every silent edit be detectable and let no downed server take the evidence with it.

For memory institutions and forensic teams: anchor the exhumation files — photos, finds, the DNA chain [4][5]; the families deserve certainty, and history, an unassailable record.

For survivors, families and oral history projects: record and seal the testimonies now, in life [1][5]; temporal priority will be their shield when the forgeries come.

For museums and photographic collections: anchor the digitizations — the pre-AI horizon closes a little more every day [2][5].

For the media: label every synthetic image of real events or people [2] and anchor your archives — your retrospective credibility is heritage too.

For AI companies and legislators: follow UNESCO's recommendations — work with survivors, educators and historians [1] — and give the Democratic Memory Law [4] the forensic and anchoring means to make its work irreversible.

A machine invented a "Holocaust by drowning" that never existed [1]. Let the world, when someone "documents" it, be able to answer with an anchored corpus that proves what is memory and what is fabrication. For those who can no longer testify, **time2act**.

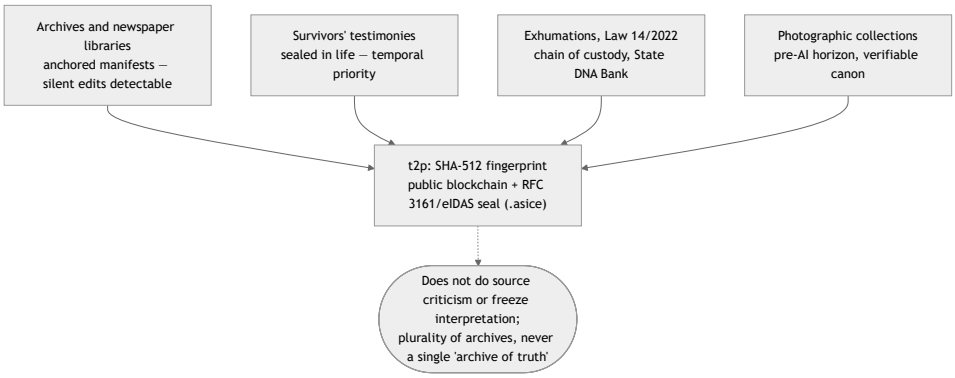


Figure 61.1 — Four fronts of one defense (SHA-512 fingerprint, RFC 3161/eIDAS timestamp, anchoring to the public blockchain, .asice): the anchored manifests of archives and newspaper collections (silent edits detectable), survivors' testimonies sealed in life (temporal priority against fabrication), exhumation files with an anchored chain of custody (Law 14/2022, State DNA Bank) and the pre-AI horizon of photographic collections (the verifiable canon against synthetic history). Limits: the seal does not do source criticism or freeze interpretation; plurality of archives, never a single "archive of truth." Sources: UNESCO/WJC; Maldita.es; Internet Archive (2024 breach); Law 14/2022.

---

[1] UNESCO — *New UNESCO report warns that Generative AI threatens Holocaust memory* (June 2024, with the World Jewish Congress; ChatGPT and the false "Holocaust by drowning"; testimony quotes fabricated by Bard; oversimplification; recommendations to technology companies and education systems); press release in Spanish (PDF). https://www.unesco.org/en/articles/new-unesco-report-warns-generative-ai-threatens-holocaust-memory · https://articles.unesco.org/sites/default/files/medias/fichiers/2024/06/PR_New_UNESCO_report_warns_that_Generative_AI_threatens_Holocau ↩

[2] Maldita.es — *El papa, Margarita Salas y escenas de la Guerra Civil generadas con IA, en portada: cómo el uso de imágenes falsas en webs y revistas puede desinformar* (February 25, 2025; synthetic images of real events and people published without a label; alteration of historical perception and distrust in the media). https://maldita.es/malditatecnologia/20250225/imagenes-IA-acontecimientos-reales-desinformar/ ↩

[3] Forbes — *Internet History Hacked, Wayback Machine Down — 31 Million Passwords Stolen* (October 2024; 7 TB leaked; GitLab token exposed since late 2022; simultaneous DDoS on the Wayback Machine); NPR — *Hackers steal information from 31 million Internet Archive users*; Hispasec — *Internet Archive sufre su segunda brecha* (more than 800,000 support tickets exposed). https://www.forbes.com/sites/daveywinder/2024/10/10/internet-hacked-wayback-machine-down-31-million-passwords-stolen/ · https://www.npr.org/2024/10/20/nx-s1-5159000/internet-archive-hack-leak-wayback-machine · https://unaaldia.hispasec.com/2024/10/internet-archive-sufre-su-segunda-brecha-de-seguridad-en-octubre-de-2024.html ↩

[4] Infobae — *La Ley de Memoria Democrática impulsa la exhumación de fosas y la reparación de 9.000 víctimas: todavía quedan más de 10.000 cuerpos por recuperar* (November 18, 2025; 8,941 remains located and recovered; only 70 identified and returned; Integrated Map of Graves: 5,848 graves, ≈1,500 opened; State DNA Bank); El Español — *La nueva Ley de Memoria entra en vigor: el Gobierno quiere exhumar otros 25.000 cadáveres* (October 2022). https://www.infobae.com/espana/2025/11/18/la-ley-de-memoria-democratica-impulsa-la-exhumacion-de-fosas-y-la-reparacion-de-9000-victimas-todavia-quedan-mas-de-10000-cuerpos-por-recuperar/ · https://www.elespanol.com/espana/20221020/nueva-ley-memoria-gobierno-quiere-exhumar-cadaveres/712178865_0.html ↩

[5] t2p — product manual (SHA-512 fingerprint at the source; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; batch manifests; content in each archive's systems, only fingerprints on the chain). See chapters 13 and 15; official statements, in chapter 86; war crimes evidence, in chapter 20. ↩

Whistleblowing: Proving Without Giving Yourself Away

"One word of truth shall outweigh the whole world."

— Aleksandr Solzhenitsyn, Nobel Lecture (1970)

In June 2017, a twenty-five-year-old NSA contractor named Reality Winner printed a classified report on Russian interference in the U.S. elections and mailed it to *The Intercept*. The newsroom published a **scan of the printed copy**. A few days later, the FBI was knocking on her door. The document betrayed her with an invisible treachery: tiny **yellow dots** that almost every color laser printer embeds on every page — the machine identification code — which, decoded with the Electronic Frontier Foundation's tool, revealed the printer model, its serial number (29535218) and **the exact date and time of printing: May 9, 2017, at 6:20** [6][8]. The NSA's internal logs showed that only six people had accessed the report, and only one had printed it. Reality Winner was sentenced to sixty-three months in prison, the longest sentence ever imposed in the United States on a journalistic source [7].

Her case contains the warning that runs through this chapter and makes it a mirror of the whole book: **anonymity is hard, and no technology can save a careless whistleblower**. But it also poses the most delicate challenge of all: how can someone prove that a document is authentic — so that power cannot dismiss it as a forgery — **without, at the same time, giving themselves away**?

The irreducible tension

Society has recognized that whistleblowers are essential and has built defenses. The **European Whistleblower Protection Directive** of 2019 turned the confidentiality of the reporter's identity into a **legal obligation**, with secure channels and protection against retaliation; in 2025, the Court of Justice of the EU fined five states some €39 million for failing to transpose it on time — Germany €34 million [2]. And the **Freedom of the Press Foundation**, the American NGO, maintains **SecureDrop**, a system for the anonymous delivery of documents over the Tor network used by more than sixty newsrooms — the *New York Times*, the *Guardian*, the *Washington Post*. Its philosophy is revealing: source protection is not only legal but **architectural** — the system is designed so that the journalist never learns the informant's identity or network data [3].

But the architecture of surveillance is just as pervasive. As the EFF showed by decoding the printer dots, the paper that comes out of a machine already "signs" who printed it and when. And it is not the only vector: **stylometry** — the fingerprint of writing style — can attribute the authorship of a text with 80 to 95% accuracy [9]. The tension, then, is irreducible: proving **what** without revealing **who**.

Let the truth survive without the source

And yet it can be done. The biggest leak in history proves it. The **Panama Papers** — 11.5 million documents, 2.6 terabytes from the firm Mossack Fonseca — were handed over by a source who identified themselves only as "John Doe," and who imposed an absolute rule: encrypted channels only, never a face-to-face meeting [4]. More than three hundred and fifty journalists from eighty countries, coordinated by the international consortium ICIJ, verified their authenticity and published the revelations, which brought down a prime minister and won a Pulitzer Prize. A decade later, John Doe's identity remains unknown. In their manifesto, they demanded for legitimate whistleblowers "immunity from government retribution, without exceptions" [5]. The lesson is powerful: verifiability and anonymity can coexist. It was possible to prove that 2.6 terabytes of documents were genuine without needing to identify who had leaked them.

The proof of existence as a shield: what is NOT sealed

This is where **t2p** offers a tool with a design inverted with respect to the entire rest of the book. The whistleblower computes the document's **SHA-512 fingerprint on their own device** and anchors *only* that fingerprint to the **public blockchain**, with an **RFC 3161 timestamp** conforming to eIDAS. This proves two things verifiable by anyone — that the document **is authentic and has not been altered**, and that **it existed before a given instant** — without uploading the document or any personal data.

And here enters this chapter's key piece: **selective disclosure**. Through zero-knowledge proofs, a journalist or a court can verify that the document they received matches the sealed fingerprint — that is, that it is genuine and predates the leak — **without discovering who sealed it**. What is proven is a property of the document, not the identity of its author. This dismantles the "false or doctored document" defense: the accused organization cannot claim the leaker fabricated it after the fact, because the fingerprint and the seal anchor it to an earlier date.

But the most important thing is what is **not** sealed. Unlike the forensic flow in the rest of the book, here the GPS, the proof of life and the device metadata are deliberately omitted — precisely what would give the source away. The **.asice** container can be handed to the journalist **offline**, without the leaker's real identity ever touching the network. It is the conceptual turn of this work: the **same** technology that in the chapters on photojournalism and human rights was a **risk** — anchoring metadata can betray a source — becomes, configured with selective disclosure and metadata omission, a **shield**. The difference is not the blockchain, but the design of what goes into it and what is proven without being revealed.

And the limit must be said in all its rawness, because here it is vital and closes the book's circle. **Anonymity is hard, and technology does not save a careless leaker.** The Reality Winner case screams it: not even the strongest encryption would have helped when she was betrayed by the printer dots, the folds in the paper and the internal access logs. A fingerprint proves existence, integrity and moment — **not** that the content is true or that it was obtained lawfully: "garbage in, garbage on-chain." And there remain deanonymization vectors t2p does not solve: network metadata (Tor or SecureDrop must be used separately), stylometry, the physical marks of the medium. The golden rule of this front is the most counterintuitive in the whole book: what protects you **is what you decide not to seal**.

The promise and the peril

The promise is precious for democracy. A tool that lets a whistleblower prove a document is authentic and predates a certain date, without revealing who they are, would give investigative journalism and the courts incontestable evidence while protecting the person who risked everything to tell the truth. It is making the truth survive without the source falling.

The peril, here, is above all that of **false security**. Selling the proof of existence as a guarantee of anonymity would be irresponsible and could cost someone their freedom: anonymity depends on a discipline — Tor, clean devices, care with the physical medium and with style — that no seal replaces. Nor does the seal make a false document true or legitimize unlawful obtainment. The proof of existence is a layer that protects *if and only if* it is designed to omit what gives you away; never a silver bullet against a surveillance machine that the Winner case showed to be implacable.

Moreover, the timestamped seal becomes a **temporal anchor** that can be exploited as an instrument of coercion. In an investigation, whoever controls the seal's time can show that a document was "sealed" at a specific moment, and that trail can serve to pinpoint the exact instant a person exposed sensitive information. Carissa Véliz has shown how this pressure translates into self-inhibition: "between one third and two thirds of writers avoided certain topics for fear of government surveillance" [11]. The immutable record, which supposedly protects the truth, can end up becoming a threat to freedom of expression: the mere knowledge that the seal leaves a record of the moment can deter legitimate whistleblowing and force potential informants to self-censor, undermining the very goal of preserving the truth.

It's time to act

So, for whistleblowers: prove the authenticity of what you reveal without anchoring what gives you away, and combine it with secure channels like SecureDrop and strict discipline. For journalists and leak platforms: offer selective disclosure so a document can be verified without knowing its source. For legislators: comply with and strengthen the whistleblower protection directive, which too many states transposed late. And for the public: remember that without those who dare to speak — and without tools that protect them — much corruption would never come to light.

Invisible yellow dots sent Reality Winner to prison; an anonymous person called John Doe shook the world and remains free. The difference, often, is what gets revealed unintentionally.

Giving the whistleblower a way to prove the truth **without giving themselves away** — sealing the document and hiding the person — is, on the front that closes this part, time2act.

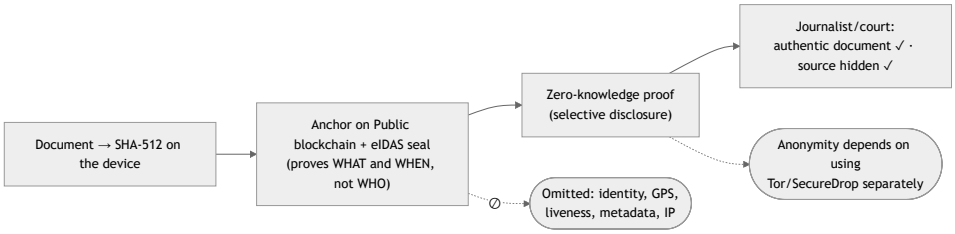


Figure 62.1 — The whistleblower computes the document's SHA-512 fingerprint on their device and anchors it to the public blockchain with an eIDAS timestamp: WHAT and WHEN are proven, not WHO. The layer that would betray them — identity, GPS, proof of life, metadata, IP — is deliberately omitted: what protects is what is not sealed. With a zero-knowledge proof, the journalist or the court verifies that the document is authentic and predates a certain date, without knowing the source. Limit: the chain does not block printer tracking dots (the Winner case), stylometry or network metadata. Sources: ICIJ (Panama Papers); EFF; Freedom of the Press Foundation.

---

[1] EUR-Lex — *Directive (EU) 2019/1937 on the protection of whistleblowers* (confidentiality of identity; secure channels; deadlines). https://eur-lex.europa.eu/eli/dir/2019/1937/oj/eng

[2] eucrim — *ECJ orders several member states to financial penalties for failing to transpose the Whistleblowers Directive* (March 2025; ~€39 million; Germany €34 million; Estonia €500,000 + €1,500/day). https://eucrim.eu/news/ecj-ordered-several-member-states-to-financial-penalties-for-failing-to-transpose-whistleblowers-directive/ ↩

[3] Freedom of the Press Foundation — *SecureDrop* (more than 60 newsrooms; Tor hidden service; architectural source protection). https://securedrop.org/overview/ ↩

[4] ICIJ — *Panama Papers* (11.5 million documents; 2.6 TB; >350 journalists; the source "John Doe"). https://www.icij.org/investigations/panama-papers/ ↩

[5] OCCRP — *John Doe: "The Revolution Will Be Digitized"* (the Panama Papers source's manifesto). https://www.occrp.org/en/investigation/panama-papers-source-john-doe-the-revolution-will-be-digitized ↩

[6] Errata Security (Robert Graham) — *How The Intercept outed Reality Winner* (printer tracking dots; serial 29535218; 9-May-2017, 6:20). https://blog.erratasec.com/2017/06/how-intercept-outed-reality-winner.html ↩

[7] Freedom of the Press Foundation — *Reality Winner sentenced to 63 months* (the longest U.S. sentence for a press source). https://freedom.press/issues/whistleblower-reality-winner-sentenced-longest-prison-sentence-history-federal-leak-cases/ ↩

[8] Electronic Frontier Foundation (EFF) — *List of printers which do or do not display tracking dots* (machine identification code). https://www.eff.org/pages/list-printers-which-do-or-do-not-display-tracking-dots ↩

[9] Wikipedia — *Stylometry / Adversarial stylometry* (authorship attribution with 80-95% accuracy). https://en.wikipedia.org/wiki/Stylometry ↩

[10] Wikipedia — *Frances Haugen* (the Facebook leak; chose to reveal her identity). https://en.wikipedia.org/wiki/Frances_Haugen

[11] Carissa Véliz — *The Ethics of Privacy and Surveillance* (Oxford UP, 2024). ↩

Research Integrity: The Datum No One Can Repeat

"The first principle is that you must not fool yourself — and you are the easiest person to fool."

— Richard Feynman, "Cargo Cult Science," Caltech (1974)

In April 2002, the physicist Lydia Sohn, then at Princeton, was comparing the plots of a young Bell Labs researcher named Jan Hendrik Schön. The measurements were presented as taken on different devices and at very different temperatures; but they showed an impossible feature: **exactly the same noise**. The noise of an electronic measurement is random, unique as a fingerprint; for two independent measurements to share the same pattern could mean only one thing. Sohn, together with Paul McEuen of Cornell, had just uncovered the greatest fraud in modern physics: Schön had copied and pasted the data. The Beasley committee confirmed misconduct on sixteen of twenty-four charges and close to twenty-nine papers were retracted, eight of them in *Science* and seven in *Nature* [6]. The year before, Schön had been publishing, on average, a new paper every eight days.

That copied noise raises this chapter's central question. It is not about who had the idea first or whether a degree is authentic — that belongs to chapter 19 — but about something more fundamental: **did this datum, this image, exist unmanipulated at the moment of acquisition?** Because science, the most reliable knowledge system we have built, rests on a chain of evidence — pixels, blots, raw files — which, like the fraud's, can be fabricated.

The bill for science that cannot be repeated

The scale of the problem is dizzying. In 2023, **more than ten thousand scientific papers** were retracted, an all-time record; the bulk came from the publisher Hindawi, a Wiley subsidiary, which withdrew more than eight thousand from special issues infiltrated by *paper mills* — companies that manufacture and sell fake papers [1]. As **Ivan Oransky and Adam Marcus**, cofounders of **Retraction Watch**, the reference database on retractions, warn, this figure is "only the tip of the iceberg": the correction system lags far behind industrialized fraud.

And what is not outright fraud often cannot be repeated either. The company Amgen managed to reproduce only **11%** of fifty-three landmark oncology studies; the Reproducibility Project: Cancer Biology, published in *eLife* in 2021 after nearly a decade of work, found that the replicated effects were, at the median, **85% smaller** than the originals [10]. On top of all this comes deliberate manipulation: the psychologist Diederik Stapel has fifty-eight papers retracted for inventing entire datasets — when asked about the schools where he had supposedly run the experiments, he admitted that **those schools did not exist** — and the individual record was held in July 2023 by the anesthesiologist Joachim Boldt, with one hundred and eighty-four retractions. When the evidence is easy to fabricate and hard to verify, the whole edifice trembles.

Fraud before and after pressing "capture"

Data fraud enters through two doors. The first is **image manipulation** after the fact. The microbiologist **Elisabeth Bik**, in her foundational study with Arturo Casadevall and Ferric Fang (*The Prevalence of Inappropriate Image Duplication in Biomedical Research Publications*, mBio, 2016), visually inspected **20,621 papers** from forty journals and found that **3.8%** contained problematic figures, half with signs of deliberate manipulation: Western blot bands duplicated, rotated, cropped and pasted [2]. And as Bik warns, since she looked at only one type of data, the real prevalence is almost certainly higher. The good news is that automated screening works: in a one-year pilot, the American Society for Microbiology ran 2,627 accepted manuscripts through the ImageTwin tool and flagged 410 with image problems, to the point of revoking acceptance for six [3]. The most notorious case was uncovered from a village in Wales by the biologist **Sholto David**, who in January 2024 published on his blog images from papers of Harvard's Dana-Farber Cancer Institute that looked cropped and pasted; the institute moved to retract six studies and correct thirty-one. Almost two years later, in December 2025, Dana-Farber paid **fifteen million dollars** to settle a federal lawsuit, admitting that images and data had been "misrepresented or duplicated" in grant applications to the national institutes of health [4].

The second door is subtler: **outcome switching**. The team of **Ben Goldacre**, at the University of Oxford, monitored with the COMPare project sixty-seven clinical trials published in the five major medical journals and discovered that, on average, only **76.3%** of the prespecified primary outcomes were correctly reported, while 365 new outcomes were added that no one had announced beforehand [8]. This is what is known as *outcome-switching*: measuring one thing, finding nothing, and reporting another as if it were what you were looking for. The defense against this is **preregistration**: depositing the protocol before starting. When the United States made it mandatory for drug trials in 2007, registration went from 64% to 100% and results reporting from 10% to 100%. And the **Center for Open Science**, driven by Brian Nosek, has extended *Registered Reports* — where the journal accepts the study on its methods, before knowing the results — from ninety-one journals in 2018 to more than three hundred in 2022. The underlying idea is luminous: if the decision to publish is made *before* seeing the results, the incentive to dress them up disappears.

The proof of existence: sealing the datum at acquisition

This is where **t2p** acts with surgical precision, because it attacks the problem at its exact origin. At the instant of **acquisition** — when the confocal microscope, the gel reader, the cytometer or the sequencer generates the raw file — the **SHA-512 fingerprint** of the original file is computed on the device itself and anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS. An unalterable record remains: "this exact pixel existed at this date and time." Any later crop, rotation, band copy or contrast adjustment produces a different fingerprint that no longer matches the one sealed at origin: the manipulation becomes **mathematically detectable and datable**. Schön's copied noise would not have survived this check.

The same applies to **preregistration with a certain date**: the protocol, the primary outcomes and the analysis plan are sealed *before* any data is collected, so that if new outcomes appear later, the date of the **.asice** container proves they were not prespecified — it is the cryptographic equivalent of what COMPare needs to catch *outcome-switching*, but verifiable without depending on a central registry. And the combination of **forensic GPS** and **proof of life** anchors the capture to a real place, moment and human operator: against Stapel's fraud, it would demand authentic coordinates and time instead of a spreadsheet invented in an office.

The limit must be said in all its rawness, because here it is the heart of the matter. "Garbage in, garbage on-chain": the seal guarantees that a file existed unaltered at a date and place, but **it does not guarantee that the datum is true or that the experiment is good**. A dishonest researcher can seal at origin a datum **fabricated beforehand** — numbers invented on the instrument itself, a fake sample, a rigged control: the fingerprint will be perfect and the provenance impeccable, and yet the science will be fraud. Nor does the seal detect honest errors or reproducibility problems — a result can be perfectly dated and be irreproducible through bad methodology — nor does it replace the expert eye of Bik or David, who catch manipulation *prior* to pressing "capture." The proof of existence moves the point of trust from "I believe your figure" to "I believe your acquisition," a real advance against later manipulation, but **it does not close off fabrication at origin**.

The promise and the peril

So the promise is immense for a system that moves trillions in funding and on which medicine and technology depend. Sealing the raw datum at acquisition and the protocol before the experiment would make it much harder to duplicate a band, reuse noise or add outcomes after the fact, and would give reviewers and integrity committees a verifiable chain of custody instead of a figure that must be taken on faith. In a science eaten away by *paper mills*, making the evidence unalterable is protecting the foundation of knowledge.

The peril calls for the usual nuance. First, as said: the seal does not make a false datum true; if the fabrication happens before capture, the chain immortalizes the lie with impeccable provenance. Second, overconfidence: no fingerprint replaces independent replication or the expert eye. And third, adoption: if sealing is optional or happens after acquisition, it loses almost all its value, and the absence of a seal does not by itself prove fraud. The proof of existence is a powerful tool against manipulation and outcome switching, but it is not a certificate of scientific truth.

It's time to act

Sealing the raw file at capture and preregistering the protocol with a certain date guarantees the provable integrity of the data.

Journals must demand the sealed raw datum and compare its fingerprint, just as is done with automated image screening.

Funders must condition grants on a verifiable chain of custody for the data, something that would have prevented the Dana-Farber case.

Institutions must protect those who, like Sholto David, uncover manipulation, and not penalize them.

A copied noise brought down the biggest fraud in physics when someone compared two fingerprints. Giving every scientific datum a fingerprint sealed at birth — without forgetting that final truth requires repeating the experiment — is, on this front of knowledge, time2act.

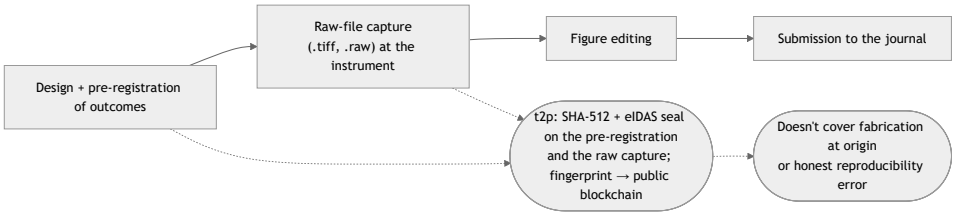


Figure 63.1 — Across the life of an experiment, fraud can enter at the design (outcome switching: COMPare found that only 76.3% of primary outcomes were reported properly), at the sample (fabrication at origin, like Stapel), at the capture of the raw file or at the editing of the figure (band duplication: Bik found it in 3.8% of 20,621 papers). The seal at acquisition — SHA-512 fingerprint + eIDAS timestamp — covers the preregistration of the protocol and the capture of the raw datum, and makes any later edit detectable; but it does not prove that a datum fabricated at origin is true. Sources: COMPare (Oxford); Bik et al. (mBio 2016); ASM/ImageTwin.

---

[1] Nature (Richard Van Noorden) — *More than 10,000 research papers were retracted in 2023 — a new record* (Hindawi/Wiley, >8,000). https://www.nature.com/articles/d41586-023-03974-8 ↩

[2] Elisabeth Bik, Arturo Casadevall and Ferric Fang — *The Prevalence of Inappropriate Image Duplication in Biomedical Research Publications* (mBio, 2016; 20,621 papers, 3.8%). https://journals.asm.org/doi/10.1128/mbio.00809-16 ↩

[3] American Society for Microbiology / mBio — *ImageTwin pilot* (2,627 manuscripts; 410 with image problems; 6 acceptances revoked). https://journals.asm.org/doi/10.1128/mbio.01990-25 ↩

[4] U.S. Department of Justice — *Dana-Farber Cancer Institute agrees to pay $15M to settle fraud allegations* (images/data misrepresented in NIH grants). https://www.justice.gov/opa/pr/dana-farber-cancer-institute-agrees-pay-15m-settle-fraud-allegations-related-scientific ↩

[5] Stanford Board of Trustees — *Report on the scientific work of Marc Tessier-Lavigne* (2023; the president's resignation). https://boardoftrustees.stanford.edu/news/board-trustees-statement-release-report-and-announcements

[6] American Physical Society — *The Schön scandal* (Beasley committee; misconduct on 16 of 24 charges). https://www.aps.org/apsnews/2022/08/september-2002-schon-scandal-report ↩

[7] Retraction Watch (Ivan Oransky and Adam Marcus) — *The new retraction record holder* (Joachim Boldt, 184 retractions, July 2023). https://retractionwatch.com/2023/07/12/the-new-retraction-record-holder-is-a-german-anesthesiologist-with-184/

[8] Ben Goldacre et al. (COMPare, University of Oxford) — *COMPare: outcome-switching in clinical trials* (Trials, 2019; 76.3% of primary outcomes properly reported). https://pmc.ncbi.nlm.nih.gov/articles/PMC6375128/ ↩

[9] Center for Open Science (Brian Nosek) — *Registered Reports* (from 91 to more than 300 journals; the protocol before the results). https://www.cos.io/initiatives/registered-reports

[10] eLife — *Reproducibility Project: Cancer Biology* (Amgen, 11% of 53 studies; replicated effects 85% smaller). https://elifesciences.org/articles/71601 ↩

Pensions: Proving You Are Still Here, Today

"Old age is honored only on condition that it defends itself, maintains its rights, is subservient to no one, and to the last breath rules over its own domain."

— Cicero, De Senectute, §38 (44 BC)

On July 28, 2010, officials in Tokyo went to Sogen Kato's home to congratulate him: according to the register, he was the oldest man in the city, at one hundred and eleven, and Respect for the Aged Day was approaching. When they opened his room, they found a mummified body in pajamas, surrounded by newspapers from 1978. Kato had been dead inside his own house for thirty-two years while his **survivor's pension** kept dropping, month after month, into the family account. The family had wrongfully collected some 9.5 million yen; his granddaughter was convicted [1]. The case exposed a national hole: the Japanese police could not confirm whether **234,354 people** listed as centenarians were still alive — 77,118 of them would have been over one hundred and twenty — because the family registers had gone decades without updating [2].

That mummified body contains this chapter's question, which is not chapter 47's — "who are you?", answered once — nor chapter 49's — "did the aid arrive?" — but a more persistent one: **are you still alive, today, and will you be again in three months?** Because a pension is not a single act but a right collected over years, and the most silent fraud is not an event but a **non-event**: a death no one declares.

The state that does not know who has died

The problem is universal and expensive. In the United States, a 2025 report by the Social Security inspector general found that the automatic rejection of death notices caused **$327 million** in payments to dead people; another report had already detected **6.5 million** Social Security number holders aged one hundred and twelve or older **with no date of death** in their records [3]. In Greece, under pressure from the troika, a census of pensioners suspended benefits to some **200,000 people** who were not entitled to them or were already dead — among them, some nine thousand listed as "centenarians" collecting retirement — saving up to **€800 million a year** [4]. And in South Africa, the grants agency SASSA admitted having paid some **140 million rand to close to 75,000 deceased beneficiaries** [5]. The root of it all is the same: a **civil registry that does not reliably know who has died**.

And there is an uncomfortable nuance that must be said: a good part of the fraud is committed not by the dead person, but by whoever controls the system. In South Africa, an analysis by the Institute for Economic Justice attributes close to 75% of grant fraud to officials and intermediaries, not to beneficiaries: insiders who create "phantom beneficiaries," duplicate cards and change bank details. That is why, as the International Social Security Association (ISSA) warns in its analysis *Proof of life in international social security* [9], proof of life must avoid depending on a central repository that can be manipulated from inside.

Proving you are alive, today and in three months

The traditional answer — making the pensioner come to a counter once a year — is humiliating for the elderly and impossible for those who live far away. Several countries have sought digital alternatives. India created the digital life certificate **Jeevan Pramaan**, which last year generated some 19 million certificates, more than eleven million of them through facial authentication, sparing the physical appearance [6]. Brazil inverted the logic with its *prova de vida*, cross-checking public databases. And Spain, with the **VIVESS** app, lets pensioners living abroad prove they are still alive without traveling.

But here a new threat appears. As the Brazilian legal journal **Conjur** warns in its article *Deepfakes, provas digitais e fraudes no INSS* [10], deepfakes threaten an "epistemic apocalypse" in which a video stops serving as proof; there are already fake agents collecting retirees' faces and voices "for the proof of life." The lesson connects with this whole book: "a credible video" is not enough; what is needed is proof that is **cryptographically verifiable**.

The recurring proof of existence: the challenge the dead cannot answer

This is where **t2p** offers a mechanism tailor-made for this problem. In each attestation cycle — monthly, quarterly or yearly — the pensioner performs a **liveness gesture** on their device: turning their head, saying a phrase, answering an active anti-spoofing challenge. The device locally computes the **SHA-512 fingerprint** of the package formed by {liveness proof + **challenge of the day** + timestamp + forensic GPS}, anchors it to the **public blockchain** with an **RFC 3161 seal** conforming to eIDAS and closes it in an **.asice** container verifiable offline. What gets recorded on the chain is not the face, but **the proof that a living human answered an unrepeatable challenge at a specific instant and place**.

The key is the **fresh challenge**, a random value of the day that acts as an antidote to reuse: a recording from six months ago does not contain today's phrase, and its fingerprint will not match. This strikes directly at the phantom-pensioner fraud — **the dead cannot answer today's challenge** — while the GPS and the timestamp detect impossible patterns, such as one device attesting for fifty beneficiaries in the same minute. And all of it **without a central biometric database**: unlike Aadhaar or the Brazilian system's databases, the biometric template **never leaves the device**; only a fingerprint lives on the chain. There is no honeypot to steal. Moreover, the container is generated offline and anchored when there is network coverage, exactly what rural or expatriate pensioners need — the ones who fail today's in-person systems.

The limit must be said in all its rawness, because here it is delicate. "Garbage in, garbage on-chain": t2p can certify unforgeably that a living human answered a unique challenge, but **it does not guarantee that this human is the legitimate beneficiary** if the initial enrollment was fraudulent or if a caregiver coerces the person. If a relative forces a living elder to perform the gesture in order to keep collecting a pension they are not entitled to, the chain will record a perfectly valid proof of life: the fingerprint is honest about what the sensor saw, not about the legal legitimacy of the payment. Second, **it does not solve the inverse problem**: as the British system's operation LEAP shows — **119,050 underpayments worth £736 million** [8], with more than fifteen thousand people who can no longer be found because they have died — the real damage is often the state that **owes money and does not pay it**. Proving you are alive does not force the administration to calculate your entitlement correctly. And third, liveness is not infallible and the technological barrier **excludes**: as the Indian outlet **The Wire** documents [7], there are elders denied their pension because their worn fingerprints do not match the biometric database — one case, that of an eighty-three-year-old woman, ended in death after sixteen months without pension or rations. That is why fingerprint authentication was eventually made **non-mandatory**. Demanding a digital gesture from someone with dementia, without a phone or without coverage reproduces the digital divide; the periodic proof must coexist with an assisted, non-punitive channel, or it becomes a machine for suspending the pensions of the weakest.

The promise and the peril

The promise is human and quantifiable. A recurring proof of life, verifiable and with no central database, would make collecting after an undeclared death almost impossible — the hole of Tokyo, of Athens, of Pretoria — without turning the elderly into a filed population or forcing them to stand in line. In systems that pay pensions to hundreds of millions of people, sealing liveness over time protects the public purse and the dignity of the recipient at once.

The peril calls for the nuances we have stated. The first, coercion and fraudulent enrollment, which no seal solves. The second, not confusing proof of life with eligibility control or with the quality of the death registry, which is the root of the problem. And the third, above all, inclusion: a tool designed against fraud cannot end up excluding precisely the most vulnerable. The proof of existence hardens the step "are they alive and present, today?", which is exactly where phantom pensioners fail, but it replaces neither inclusion policy nor a civil registry that knows who has died.

It's time to act

For **pension funds and social security**: adopt a recurring proof of life with a fresh challenge and no central biometric database, so the dead cannot collect and the living do not have to travel. For **civil registries**: close the hole at the source — knowing reliably who has died — which is the root cause. For **regulators**: require an assisted, non-punitive channel for those who cannot perform the gesture, and watch that proof of life does not become exclusion. And for **pensioners themselves and their families**: a tool that proves they are still alive without handing their biometrics to any server.

A man kept collecting his pension thirty-two years after dying because no one made him prove he was alive. Giving every pensioner a way to demonstrate, again and again, that they are still here — with a challenge the dead cannot answer and without filing the living — is, on this front that touches everyone's old age, time2act.

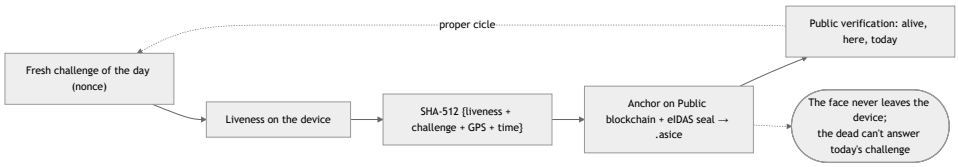


Figure 64.1 — Each attestation cycle starts from a fresh challenge of the day (nonce); the pensioner performs the liveness gesture on the device; the SHA-512 fingerprint of {liveness + challenge + GPS + time} is computed locally, anchored to the public blockchain with an eIDAS seal and closed in an .asice verifiable offline. The unrepeatable challenge prevents reusing old recordings: the dead cannot answer today's challenge. The face never leaves the device; only a fingerprint travels on the chain. The cost of not knowing who is still alive: Japan, 234,354 unconfirmed centenarians; U.S., $327 million to the deceased; South Africa, 140 million rand to 75,000 dead. Sources: France 24; SSA OIG; SASSA/Moneyweb.

---

[1] Wikipedia (with BBC / AFP / Japan Times) — *Sogen Kato* (mummified 32 years; survivor's pension collected; ~9.5 million yen). https://en.wikipedia.org/wiki/Sogen_Kato ↩

[2] France 24 — *Over 230,000 listed centenarians missing in Japan* (234,354 unconfirmed; 77,118 aged 120+). https://www.france24.com/en/20100910-over-230000-listed-centenarians-missing-presumed-dead-government-report-japan ↩

[3] SSA Office of the Inspector General — *State death report discrepancies led to $327M in improper payments* (2025) and *6.5 million holders aged ≥112 with no date of death in the Numident* (2015). https://oig.ssa.gov/news-releases/2025-03-19-state-death-report-discrepancies-led-to-327m-in-improper-payments-and-cost-ssa-millions-of-dollars-in-unnecessary-workloads/ ↩

[4] Reuters (via HuffPost) — *Greece stops pensions and benefits for 200,000 people* (~9,000 "centenarians"; savings of up to €800 million/year). https://www.huffpost.com/entry/greece-stops-pensions-benefits-200000-people-fraudulently_n_1451423 ↩

[5] Moneyweb — *SASSA explains why it paid R140m in grants to 75,000 dead people* (2025; 75% of the fraud by officials, according to the IEJ). https://www.moneyweb.co.za/news/south-africa/sassa-explains-why-it-paid-r140m-in-grants-to-75-000-dead-people/ ↩

[6] Outlook Money — *Jeevan Pramaan: digital life certificate* (~19 million certificates; ~11.5M via facial authentication). https://www.outlookmoney.com/retirement/digital-life-certificate-over-14-million-super-seniors-generated-online-jeevan-pramaan-patra-in-fy25-26 ↩

[7] The Wire (India) — *Digital Exclusion: Poor, Elderly Face the Brunt of Aadhaar-Based Authentication Errors* (worn fingerprints; pensions denied; fingerprint made non-mandatory). https://m.thewire.in/article/rights/digital-exclusion-poor-elderly-face-the-brunt-of-aadhaar-based-authentication-errors ↩

[8] GOV.UK (DWP) — *State Pension underpayments (operation LEAP)* (119,050 underpayments, £736 million; >15,000 already dead). https://www.gov.uk/government/publications/state-pension-underpayments-progress-on-cases-reviewed-to-30-september-2024/state-pension-underpayments-progress-on-cases-reviewed-to-30-september-2024 ↩

[9] International Social Security Association (ISSA) — *Proof of life in international social security* and the VIVESS good practice (Spanish pensioners abroad). https://www.issa.int/analysis/proof-life-international-social-security-progress-and-persistent-challenges ↩

[10] Consultor Jurídico (Conjur, Brazil) — *Deepfakes, provas digitais e fraudes no INSS: o colapso da confiança* (2025). https://www.conjur.com.br/2025-mai-28/deepfakes-provas-digitais-e-fraudes-no-inss-colapso-da-confianca/ ↩

Proof of Life: Virtual Kidnappings, the Missing, and the Voice That Wasn't

"It was my daughter's voice. It was her cries. It was her sobs. It was the way she spoke."

— Jennifer DeStefano, testimony before the U.S. Senate (13-6-2023)

Arizona, January 20, 2023, around five in the afternoon. Jennifer DeStefano answers the phone and hears the voice of her fifteen-year-old daughter — who is away at a sports competition — crying and begging for help. A man takes the line: he demands a million dollars, then lowers the ransom to fifty thousand and starts dictating transfer instructions. She does not believe the people around her telling her it could be a scam: "the voice was so real, and her crying, so real." The "kidnapping" collapsed all at once when the daughter, bewildered, called to say she was fine: she had never been in danger. There was no kidnapper — only a voice cloned with artificial intelligence. DeStefano ended up testifying before a U.S. Senate subcommittee so that no other family would have to go through that hell [1].

This chapter deals with proof of life when someone attacks it. The routine life certificate — proving every year that you are still alive to collect your pension — is in chapter 64; voice cloning technology, in chapter 68; digital replicas, in chapter 81. Here the question is the worst-possible-day question: how do you know the person on the other end is real, alive and now? The old answer — the photo with the day's newspaper — is dead: anyone can fabricate it. The new answer is cryptographic: an unpredictable challenge, answered in the moment and sealed with **t2p** — the SHA-512 fingerprint on the device, the RFC 3161/eIDAS timestamp, the anchor to the public blockchain inside an .asice container [4]. A voice can be cloned; the present cannot.

The kidnapping that does not exist

The mechanics of virtual kidnapping, as the FBI describes it, is a perfectly calibrated panic machine: a call — often from an international number — claims a relative has been kidnapped; screams are mixed in, or a voice that sounds like the loved one; an immediate ransom is demanded, by wire transfer only; and the victim is kept from hanging up or trying to contact the supposed kidnapped person. The whole deception depends on speed, because there is no kidnapping: all it takes is for the family to pay before the lie collapses under its own weight [2]. The official recommendations run against instinct: stay calm, buy time, try to reach the relative by another route, and ask questions only the real person would know how to answer [2].

And in 2025 came the next step. The alert from IC3 — the FBI's cybercrime complaint center — documents criminals using altered "proof of life": photos and videos manipulated or generated with AI from public images of the victim, fabricated to make the nonexistent kidnapping plausible. The inconsistencies the FBI recommends looking for in them — a missing tattoo, a scar that is not there — are exactly the kind of detail that every generation of synthesis erases a little better [3].

Read that carefully, because it is the core of the chapter: the virtual kidnapper's trade now includes fabricating the proof of life. Crime's oldest evidentiary ritual — "prove to me that they're alive" — has become a field of deepfakes, and the family, at the moment of maximum panic of their lives, has to play forensic expert in thirty seconds. No family should play this game by the old rules.

The counterproof: challenge, freshness and seal

The structural answer is built on an asymmetry that plays in the family's favor: the cloner can imitate a voice [1], but can only say what its operator knows; and it can fabricate an image [3], but cannot answer the unpredictable. Hence the counterproof in three pieces [4].

The first is **the challenge**: a question only the real person can answer — the shared memory, the family safe word — the same logic the FBI already recommends [2] and that we saw in the family verification of chapter 85. Personal enough that no algorithm can anticipate it; simple enough that the real person answers it without thinking.

The second is **freshness**: the answer must be from now — a short video or audio clip that includes the just-posed challenge, not recycled or prepared material. The content of the present cannot have been

synthesized in the past.

And the third is **the seal**: the answer is fingerprinted and anchored to the minute — SHA-512, RFC 3161/eIDAS timestamp, public blockchain [4] — so it can neither have been fabricated beforehand nor altered afterward.

Against this triple wall, the cloner is lost: they can imitate the timbre [1], but do not know the challenge; they can generate a convincing image [3], but not the freshness of the moment; and even if they had everything, they cannot fabricate a cryptographic seal of the present dated yesterday. The family, for its part, has a protocol instead of a panic: faced with any ransom call, demand the sealed proof of life with a challenge; if it does not come, it is the scam the FBI describes [2][3]. It is the same principle of challenge and freshness that chapter 64 applies to the routine life certificate of pensions — carried into the darkest minute.

The real hostage and the missing person

There are two scenarios graver than the scam, and both need the same instrument.

The first is the real hostage. When a kidnapping truly exists, the proof of life is negotiating currency — and its authenticity, a matter of life and death in the decisions of families and negotiators. In the deepfake era, doubt cuts both ways: an authentic video can be dismissed as fake, and a fabricated one can keep a negotiation alive for a person who is already gone. The seal's contribution is dating: a video that responds to information that only existed after a certain date — and that arrives with an anchored fingerprint [4] — attests, at minimum, that the person was alive after that date. It does not resolve the negotiation; it clears the fog that poisons it. And the same logic serves those who verify from outside — the journalism and organizations of chapter 20: dated, intact material instead of uncheckable fragments.

The second is the missing person. In Spain, the National Center for Missing Persons — a pioneer in Europe — coordinates the response to thousands of disappearances every year: 16,147 people investigated in 2024 (6% more than the year before, with 72% of cases resolved in under a week) and 16,024 in 2025, with a resolution rate of 90.7%; that year the CNDES published 928 alerts, 132 of which remained active [5]. Behind every active alert there is a family searching for a trail.

The contribution of the proof of existence here is delicate, and it must be said with the limit upfront: voluntary or nothing. The person who, by their own decision, keeps a sealed trail of moments and places on their device [4] can give an eventual investigation the last verifiable point on the map; and the "return" — real, or simulated with the absent person's voice and face [3] — can be checked with challenges and seals. An imposed trail, by contrast, would be surveillance, not protection: that border is not crossed.

The proof of existence applied

The phone rings with the worst possible news. For the first time, the family has something to do other than believe or pay.

1. The family protocol against virtual kidnapping

The family agrees on the protocol before anything happens: safe word, callback to the usual number [2] and, faced with any ransom demand, a single request — sealed proof of life with a challenge [4]. The scammer's call crashes against three walls: they do not know the challenge, they do not control the real person, and they cannot fabricate the seal of the present [3]. In the Arizona case [1], thirty seconds of protocol would have replaced the minutes of terror: the daughter, reached on the other line, would have answered the challenge from the locker room at the competition — sealed, dated, unambiguous.

2. The negotiation with dating

In a real kidnapping, negotiators demand datable proofs of life from the captors: the answer to a question posed today, sealed today [4]. Material that arrives with an attested fingerprint and freshness allows decisions — pay, wait, act — on facts and not on hopes; and material without any anchor is treated with the suspicion it deserves, because the FBI has already documented the industry of fabricated proof of life [3]. Dating does not soften the drama; it clears its fog.

3. The last point on the map

A person who freely chooses it — the solo hiker, the professional who travels to risk zones, the teenager who wants to reassure their parents without feeling watched — keeps a minimal sealed trail: chosen moments, fingerprinted on their device, with no central registry and no real-time tracking [4]. If they ever disappear, the family can offer the investigation [5] a verifiable chronology of the last normality: not where they are now — the seal does not know that, nor should it — but where and when they last were, with a certain date. In a search, the first hours and the last reliable point are almost everything.

4. The return that must be believed

The return needs proof too. Whoever reappears after a disappearance — and whoever reappears only digitally: the voice on the phone, the message "it's me, I'm fine" — can prove continuity by answering challenges and sealing the answers on the spot [4]; an impersonation built from the absent person's voice and face [3] can do neither. For families, it is the difference between well-founded hope and the cruelty of a deception that can stretch on for months; for investigations [5], a filter against the false leads that consume resources and hopes.

The limits, said plainly: the seal finds no one, frees no one and replaces neither the police nor professional negotiators; the sealed proof of life requires the person to have access to a device and a minimum of freedom to answer — conditions a real kidnapping often denies; the voluntary trail can never become an obligation or surveillance; and no protocol spares a family the fear — it only gives them, in the middle of the fear, a fact to hold on to.

The promise and the peril

The promise: families with a protocol instead of a panic [1][2], negotiations on dated facts [3], investigations with reliable last points [5] and returns that can be believed. The proof of life, brought back to life.

The perils, with the same gravity.

The first is **the imposed trail**. From the voluntary "last point on the map" to the permanent geolocation of children, partners or employees runs a slope this book refuses to go down. The proof of life is an act of the person, not upon the person: their own device, chosen moments, full revocability, no central registry [4]. Anything resembling a digital collar is the exact opposite of this chapter.

The second is **coercion over the protocol**. A real kidnapper can force the victim to answer challenges and seal answers. The seal attests to freshness, not freedom: an authentic proof of life can be the product of a gun. Serious protocols know this — the dating informs the negotiation, it does not replace it, and no family should handle this without professionals [2].

The third is **false security**. No protocol covers every case: the person without a device, the forgotten challenge, the panic that erases everything. Selling the sealed proof of life as a guarantee would be cruel; it is a tool that improves the odds, not a talisman.

And the fourth is the most delicate: technifying fear. Behind every active alert [5] there are vulnerable families — and a potential market for false technological hopes. The first response to a disappearance is the police, the public protocols and support for the families; the seal is evidentiary infrastructure at their service, never a product for selling them fear.

It's time to act

The question "are they alive?" deserves a better answer than a forgeable photo.

For families: agree on the protocol before it is needed — safe word, callback, sealed proof of life with a challenge [1][2][4]; and if the call comes, buy time: the scam lives on haste [2].

For police and negotiators: build cryptographic dating into proof-of-life protocols [3][4]; material without attested freshness deserves suspicion by default.

For the CNDES and missing persons units: accept and weigh the voluntary sealed chronologies families provide [4][5]; the last reliable point shortens searches.

For platforms and carriers: make the verification channel easy — the one-tap sealed challenge — and cooperate with official alerts [2][5].

For you: if you travel alone, if you expose yourself, if you want your loved ones to be able to know without watching you — the minimal, voluntary, sealed trail is yours [4]; and talk about the protocol at home: today's uncomfortable conversation is tomorrow's calm.

For legislators: recognize the sealed proof of life in official protocols and armor it against any drift into surveillance — voluntary, revocable and with no central registries [4].

Jennifer DeStefano heard her daughter's exact voice pleading, and the voice was lying [1]. Let the next family that receives that call have a question ready and a seal to demand — and let the next truth, of the return or of the last location, arrive with a certain date. For the oldest question — "are they alive?" — **time2act**.

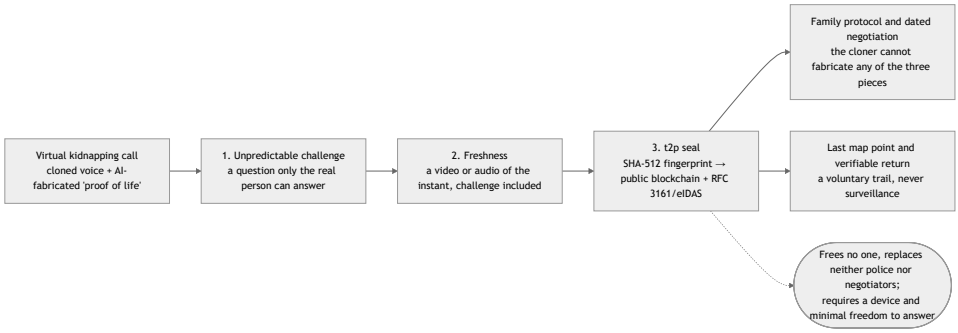


Figure 65.1 — Proof of life in the deepfake era: against virtual kidnapping with a cloned voice and AI-fabricated "proof of life" (FBI/IC3 alert, 2025), the three-piece protocol — unpredictable challenge, freshness of the moment and the t2p seal (SHA-512 fingerprint on the device, RFC 3161/eIDAS, public blockchain, .asice). Applications: the family protocol, the negotiation with dating, the voluntary last point on the map and the verifiable return (CNDES: 132 active alerts in 2025). Limits: the seal attests to freshness, not freedom; voluntary or nothing; never a substitute for police and negotiators. Sources: CNN (DeStefano case); FBI/IC3; Ministry of the Interior/CNDES.

---

[1] CNN en Español — *Una mujer cree que estafadores clonaron la voz de su hija en un falso secuestro* (the Jennifer DeStefano case, Arizona; call of January 20 around 16:55; 15-year-old daughter at a competition; $1 million ransom lowered to $50,000; "the voice was so real"; later testimony before a U.S. Senate subcommittee); La Opinión — *Utilizan voz recreada por inteligencia artificial para fingir el secuestro de una niña y pedir rescate*. https://cnnespanol.cnn.com/2023/04/29/una-mujer-cree-que-los-estafadores-clonaron-la-voz-de-su-hija-en-un-falso-secuestro-trax/ · https://laopinion.com/2023/04/29/utilizan-voz-recreada-por-inteligencia-artificial-para-fingir-el-secuestro-de-una-nina-y-pedir-rescate/ ↩

[2] FBI — *'Virtual Kidnapping' Extortion Calls on the Rise* (the scam's mechanics; international numbers; blocking contact; ransom by wire; recommendations: calm, contact by another route, questions only the real person can answer); FBI Albuquerque — *Be on the Lookout for Virtual Kidnapping Scams*. https://www.fbi.gov/contact-us/field-offices/sanantonio/news/press-releases/virtual-kidnapping-extortion-calls-on-the-rise · https://www.fbi.gov/contact-us/field-offices/albuquerque/news/press-releases/fbi-media-alert-be-on-the-lookout-for-virtual-kidnapping-scams ↩

[3] FBI / IC3 — *Criminals Using Altered Proof-of-Life Media to Extort Victims in Virtual Kidnapping for Ransom Scams* (2025 alert: photos and videos altered or AI-generated as false "proof of life," often from public images; inconsistencies such as missing tattoos or scars; report to ic3.gov with the material received); Axios — *AI-generated videos, photos used in new virtual kidnapping scams, FBI warns* (December 5, 2025). https://www.fbi.gov/investigate/cyber/alerts/2025/criminals-using-altered-proof-of-life-media-to-extort-victims-in-virtual-kidnapping-for-ransom-scams · https://www.axios.com/2025/12/05/ai-videos-photos-scam-virtual-fbi-victim-ransom-kidnapping ↩

[4] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; proof of life with challenge and freshness — an unpredictable nonce answered and sealed to the minute; voluntary trail with chosen moments; no central registry; revocability). See chapters 13, 15 and 64; the cloned voice, in chapter 68. ↩

[5] La Moncloa / Ministry of the Interior — *Las Fuerzas de Seguridad investigaron la desaparición de 16.147 personas en 2024* (+6%; 72% of cases resolved in under a week; CNDES, a pioneer in Europe); Ministry of the Interior — *Las Fuerzas de Seguridad investigaron la desaparición de 16.024 personas en 2025 y esclarecieron el 91 por ciento de las denuncias* (25,086 reports; 928 alerts published, 132 active; 11,461 unique episodes, an average of 3 reports for repeat cases). https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/interior/paginas/2025/170325-marlaska-plan-personas-desaparecidas.aspx · https://www.interior.gob.es/opencms/es/detalle/articulo/Las-Fuerzas-de-Seguridad-investigaron-la-desaparicion-de-16.024-personas-en-2025-y-esclarecieron-el-91-por-ciento-de-las-denuncias/ ↩

Sporting Integrity: The Sample Someone Swapped in the Night

"For when the One Great Scorer comes to write against your name, He marks — not that you won or lost — but how you played the Game."

— Grantland Rice, "Alumnus Football" (1908)

In the anti-doping laboratory of the Sochi Winter Olympics, one February night in 2014, an agent of the Russian secret service received through a fist-sized hole — a "mouse hole" drilled in the wall — the bottles of "dirty" urine from Russian athletes. In the next room, the team of the laboratory director, Grigory Rodchenkov, opened the supposedly tamper-proof Berlinger bottles, emptied out the doped urine and replaced it with clean urine from the same athlete, frozen months earlier, adjusting its density with **table salt**. That deception worked because the guardian of integrity — the state, the accredited laboratory, the secret service — was the one corrupting it. Months later, the London laboratory would discover that eight Sochi samples had a salt content **physiologically impossible** in a healthy human being, and that two samples from female hockey players contained **male DNA** [4].

That hole in the wall contains this chapter's question: how do you prove that a biological sample, a series of values or the result of a match has not been manipulated, when sometimes the very party who should guard them is the one altering them? Here we are not sealing a document, as in chapters 20 and 29, but **the physical contents of a body and their chain of custody**.

When the guardian is the corruptor

The Russian scandal was of industrial scale. The Canadian jurist **Richard McLaren**, appointed by the World Anti-Doping Agency (WADA) as Independent Person, documented in his *McLaren Report* (2016) the "disappearing positive methodology": at least **643 positive samples** reported as negative between late 2011 and August 2015, a figure he himself called a minimum. The report's second part concluded that **more than a thousand Russian athletes**, across more than thirty sports, had benefited from the cover-up [1][2]. And the most uncomfortable detail is that the famous "tamper-proof" bottles were not: the secret service figured out how to open their caps without breaking the visible seal, and a later investigation by the magazine *Republik* showed the redesign did not hold up either — three out of twenty-four kits could be unscrewed by hand after three days in the fridge, and one specialist found "not a single insurmountable security measure" [5]. The coverage of the McLaren Report revealed that forty-four samples, twelve of them from medalists, bore marks of tampering on their caps [3].

The lesson is brutal and is the heart of this front: **a physical seal, however sophisticated, is worth nothing if whoever controls it is complicit**. The chain of custody — which WADA defines as the sequence of parties responsible for a sample from collection to arrival at the laboratory — can be broken from inside.

From the laboratory to the betting market

Doping is not the only fraud eating away at sport; there is also match-fixing, detected indirectly by the traces it leaves in betting markets. Europol's operation VETO already identified more than 380 suspicious matches in Europe and 425 people involved in a transnational organized crime network [8]. Sportradar, in its 2024 integrity report, monitored more than 850,000 matches across seventy sports and detected 1,108 suspicious cases, half of them thanks to account betting data [9]. The proof, according to Sportradar, is no longer the confession or the visible bribe, but the anomalous movement of the odds. The biggest case starred Grigor Sargsyan, "the Maestro," sentenced to five years in Belgium in 2023 for running a network of about a hundred and eighty tennis players from more than thirty-five countries, whom he paid some seven million euros in bribes via mobile [10].

This indirect logic also permeates modern anti-doping: the Athlete Biological Passport does not "see" the banned substance, but tracks biomarkers over time and individualizes their thresholds. The cyclist Franco Pellizotti was sanctioned with two years in 2011 without a single positive test, solely on anomalous passport values — the first case upheld by the Court of Arbitration for Sport on that basis [7]. In both cases, the fraud is proven by a pattern, not by being caught in the act.

The proof of existence: sealing the sample and dating the evidence

This is where **t2p** reinforces the weakest link. At the moment of collection, the doping control officer and the athlete sign with **proof of life** a collection report that captures the SHA-512 fingerprint of the form, the **kit number** as read, the **forensic GPS** of the room, the RFC 3161/eIDAS timestamp and the athlete's face; everything is anchored to the **public blockchain** and closed in an **.asice** container. From here, every **transfer of custody** — from athlete to officer, to courier, to laboratory — generates a sealed event with time and place. Any opening, substitution or "mouse hole" would break the continuity of the anchored chain, which **cannot be rewritten retroactively** from the laboratory itself: the bank of frozen clean urine or the nighttime swap would leave a detectable temporal and geographic discontinuity. The same applies to the biological passport, where each reading can be anchored with a certain date to armor the series against the later insertion of "clean" values. And in match-fixing, t2p seals no sample but **the evidence**: Sportradar's alerts, the integrity officers' communications and the referees' decisions are sealed so that the proof presented before the Court of Arbitration or a criminal judge is demonstrably prior to any leak.

The limit must be said in all its rawness, because here it is especially strong. "Garbage in, garbage on-chain": the seal guarantees that the sample sealed at minute zero is the same one that reaches the laboratory, but **it cannot certify that its contents are the athlete's own, unmanipulated urine *before* the sealing**. If the athlete urinates with someone else's urine already, or if the doping system operates in the collection act itself with the officer's complicity, the unalterable seal will only certify a well-guarded fraud. Nor does it solve **physiologically undetectable doping** — microdosing, substances without a marker: the biological passport is indirect and probabilistic, and the chain of custody adds no analytical power to it. In match-fixing, it seals the suspicion but does not prove the intent: an anchored anomalous betting pattern does not by itself prove there was a bribe; that still demands the criminal investigation. And, above all, the Russian case teaches the hardest lesson: **when the custodian is the corruptor — the state, the accredited laboratory — the problem is not technological but one of institutional capture**, and a cryptographic seal only displaces trust, it does not eliminate it, if whoever signs the sealing act is complicit.

The promise and the peril

The promise is real and valuable. A sealed chain of custody, with time and place at every transfer, would have made Sochi's nighttime swap much harder, because the discontinuity would have been recorded in a place the laboratory cannot rewrite. And dating match-fixing evidence would give sports tribunals proof that is incontestably prior. In a world where fraud is proven by patterns, making the evidence unalterable and dated reinforces the whole edifice of integrity.

The peril calls for the nuances we have stated. The first, substitution prior to sealing, which no seal solves. The second, institutional capture: no technology saves a system where the guardian itself is the cheat, which is why the **independence of the custodian** remains irreplaceable. And the third, not confusing sealing the evidence with proving the intent. The proof of existence hardens the chain of custody and evidentiary priority; it replaces neither institutional independence nor the criminal investigation.

It's time to act

For the **anti-doping agencies** (WADA, the laboratories): seal the collection act and every custody transfer with time and place, so that no "mouse hole" can break the chain without leaving a trace. For **federations and integrity units**: date the match-fixing alerts so the evidence precedes any leak. For **regulators and sports tribunals**: demand real independence of the custodian, since no seal is worth anything in complicit hands. And for **clean athletes**: a chain of custody that protects them from the cheat and from the system that sheltered him alike.

A sample changed hands in the night through a hole in the wall because no one could prove what had happened between the seal and the analysis. Giving every sample and every piece of evidence a sealed fingerprint with time and place — never forgetting that no seal saves a corrupt guardian — is, on this front that defends fair play, time2act.

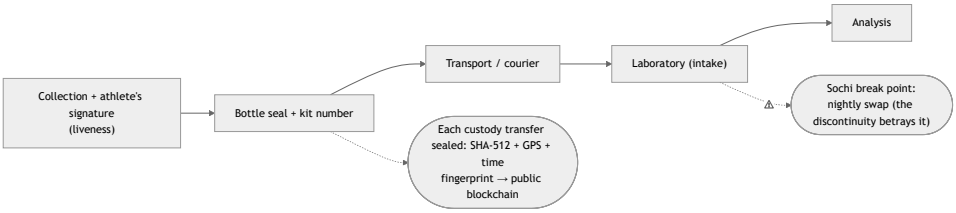


Figure 66.1 — The life of an anti-doping sample passes through five nodes: collection and the athlete's signature, sealing of bottles A and B, transport, reception at the laboratory and analysis. In Sochi, the fraud entered at the laboratory node (the "mouse hole" through which the secret service swapped urine at night). The t2p seal — forensic GPS, eIDAS timestamp and the kit number's fingerprint anchored to the public blockchain — makes that swap leave a detectable temporal and geographic discontinuity. The forensic evidence that exposed the fraud despite the "tamper-proof" bottles: 44 caps with tampering marks (12 from medalists), 8 samples with impossible salt, 2 with male DNA, ≥643 "disappeared" positives. Limit: the seal does not prove what was in the bottle before it was sealed. Sources: McLaren Report; CNN; Chemistry World.

---

[1] Wikipedia (synthesis of the *McLaren Report*) — *McLaren Report* (≥643 "disappeared" positives; male DNA in female samples). https://en.wikipedia.org/wiki/McLaren_Report ↩

[2] World Anti-Doping Agency — *Independent investigation confirms Russian state manipulation of doping control* (McLaren, Independent Person; >1,000 athletes). https://www.wada-ama.org/en/news/wada-statement-independent-investigation-confirms-russian-state-manipulation-doping-control ↩

[3] CNN — *WADA McLaren report: tampered samples* (44 samples with marks on the caps, 12 from medalists; Part 2, 2016). https://www.cnn.com/2016/12/09/sport/wada-russia-doping-mclaren-report ↩

[4] Chemistry World — *Anti-doping scientists expose cheating Russian athletes* (8 samples with physiologically impossible salt). https://www.chemistryworld.com/news/anti-doping-scientists-expose-cheating-russian-athletes/2500159.article ↩

[5] Republik — *Bad bottles* (the "tamper-proof" Berlinger bottles, breached). https://www.republik.ch/2018/02/21/bad-bottles ↩

[6] World Anti-Doping Agency — *Athlete Biological Passport* (indirect detection; hematological module 2009, steroidal 2014). https://www.wada-ama.org/en/athlete-biological-passport

[7] ESPN — *Pellizotti banned in landmark biological-passport ruling* (2 years without a positive; first case upheld by CAS). https://www.espn.com/olympics/cycling/news/story?id=6192949 ↩

[8] Europol — *Results from the largest football match-fixing investigation in Europe (Operation VETO)* (>380 matches; 425 people; >15 countries). https://www.europol.europa.eu/media-press/newsroom/news/update-results-largest-football-match-fixing-investigation-in-europe ↩

[9] Sportradar — *Betting Corruption and Match-Fixing 2024* (>850,000 matches; 1,108 suspicious; the UFDS system). https://sportradar.com/content-hub/news/sportradar-reports-notable-decline-in-match-fixing-in-2024/ ↩

[10] The Washington Post — *Grigor Sargsyan, "the Maestro," and tennis match-fixing* (5 years; ~180 tennis players from >35 countries; ~€7 million). https://www.washingtonpost.com/world/interactive/2023/tennis-match-fixing-itf-grigor-sargsyan/ ↩

Music and Live Events: The Real Human Behind the Play

"The nightingale's song is delightful because the nightingale herself gives it forth."

— John Philip Sousa, "The Menace of Mechanical Music" (1906)

In late 2018, from Cornelius, North Carolina, a middle-aged musician named Michael Smith wrote a revealing email to his two partners: "we need a TON of songs fast to get around the anti-fraud policies they're using now." He was not chasing a hit; he was chasing **statistical noise**. In his digital basement, one thousand and forty bot accounts — soon there would be ten thousand — played six hundred and thirty-six songs a day each, without rest: tracks with nonsense names and invented artists, mass-produced by an AI music company. Half a cent per stream, **661,440 streams a day**. No one was listening. For seven years no one noticed, until The MLC, the organization that administers mechanical rights in the United States, spotted the pattern. In March 2026, Smith pleaded guilty; the haul exceeded ten million dollars [1][2].

That manufactured noise contains this chapter's question. It is not chapter 40's — "did a human or an AI create it?" — because here the songs were clearly AI-made and it does not matter: the question is another one, **did a real human actually listen to this, or actually buy this ticket?** It is the fraud of *human presence* in an act of consumption.

The fraud that does not create money, it distributes it

What makes this fraud special is how it does harm. Streaming music distributes a **fixed-sum pool**: in 2025, according to the international recording federation **IFPI**, it generated more than $22 billion [6]. Each fake stream invents no new money; **it steals its share from the rest of the artists** who do have an audience. The detection company **Beatdapp** estimates that streaming fraud siphons off some two billion dollars a year — at least 10% of all activity — although, as an interested party, the figure is best taken as an industry estimate with a wide margin [3]. The platforms' response has been thresholds: Spotify decided that a track must exceed a thousand annual streams to enter the distribution, precisely to frustrate Smith's tactic of diluting a few streams across hundreds of thousands of junk tracks [4].

And there is a second identity problem, subtler still. The MLC received **$424 million** in historical mechanical royalties that were **unmatched** — the "black box": money generated by plays whose rightful owner could not be determined [5]. The question, then, is not only "was there a real listen?" but "whose is the song that was listened to?"

From the fake play to the phantom ticket

The same human-presence fraud eats away at the ticketing world. In the presale for Taylor Swift's *Eras* tour, on November 15, 2022, Ticketmaster's website received **3.5 billion requests** — four times the previous record — with some twelve million "entities," many of them resale bots, overloading the system to the point of collapse [10]. This is not an accident but a business model. The U.S. Federal Trade Commission (**FTC**) had already taken to court, under the 2016 BOTS Act, resellers who had bought more than a hundred and fifty thousand tickets with automated software and hundreds of fictitious accounts [8]. As we have seen, in 2025 devastating evidence was filed: five brokers controlled 6,345 accounts and hoarded **246,407 tickets** for 2,594 events, managed from a single screen; an executive admitted by email that the company "turned a blind eye as a matter of policy" [9]. The "four tickets per person" limit the fan saw was, in practice, fiction.

The proof of existence: the real human behind the play and the ticket

This is where **t2p** attacks the void, certifying not who created the work, but that **a real human, on a real device, at a real moment and place, performed a real action**. On the consumption side: on the listener's device, at the instant of playback, a SHA-512 fingerprint of a listening event is generated — track, device, **proof of life**, forensic GPS and the duration actually played — anchored to the **public blockchain** and sealed with the eIDAS timestamp; the player requires a liveness signal for the listen to

count. On the ticketing side: at the moment of purchase and, again, at entry to the venue, an **.asice** container is sealed binding the identity verified with proof of life, the ticket, the venue's GPS and the time of physical entry.

The effect is direct. A listen **without proof of life does not enter the count** that distributes the royalties, which sinks Smith's model: ten thousand bot accounts stop getting paid. Binding every purchase to an identity with liveness **cancels the hundreds of accounts per reseller** documented by the FTC. And the GPS-and-time seal at the door proves that whoever enters is whoever bought, not an automated bulk buyer. Human presence becomes the unit of account again.

It must be said in all its rawness, because here it is nuanced. "Garbage in, garbage on-chain": **t2p** can prove that a living human pressed "play" or crossed the gate, but **it cannot prove that the listen was genuine or honest**. If instead of bots you hire **real humans** — each passing the proof of life — to listen to junk tracks, the seal will certify them as "real humans" and the fraud will persist: the proof of life defeats bots, not human farms or paid streaming. Nor does it judge merit: a thirty-second human listen to AI noise, made only to collect, is sealed as valid. The "black box" of the $424 million is not solved by a seal either; the event existed, but no one knows **whose** the composition is, and that demands clean metadata, not just a seal. And, above all, the seal only works if the platforms integrate and honor it: the FTC's case file shows the dominant intermediary can "turn a blind eye as a matter of policy," and no cryptography corrects the economic incentive of the gatekeeper itself. In legitimate resale, moreover, binding the ticket to an identity does not prevent someone from reselling or giving their ticket to another real person.

The promise and the peril

The promise is quantifiable in the billions. A human-presence layer over streaming and ticketing would make it much harder to drain the royalty pool with bots, would cancel the resellers' armies of fictitious accounts and would return the money to artists with real audiences and the tickets to fans. In a market where fraud does not create money but distributes it badly, proving the real human is restoring the fairness of the distribution.

The peril requires the nuances we have stated. The first, human farms and incentivized streaming, since the proof of life does not stop them. The second, not confusing presence with legitimacy or with the identity of the work. The third, platform capture: no seal is worth anything if the gatekeeper decides not to honor it. And the fourth, privacy and the fan's right to resell or give away their ticket, which the anti-bot mechanics must respect. The proof of existence is a powerful tool for restoring human presence, not a silver bullet against all the sector's frauds or a substitute for fixing its incentives.

It's time to act

For the platforms (Spotify, Apple Music), the proof of human presence must be integrated into the royalty count, so that a listen without a human stops being billed. For the rights bodies (IFPI, The MLC), the seal must be required and the metadata cleaned to empty the "black box." For the ticketing industry (Live Nation/Ticketmaster) and the regulators (the FTC), purchase and entry must be bound to a real identity — and the blind eye must end. And for artists and fans, the system must guarantee that both the listen and the ticket belong to humans, not bots.

A bot farm got paid for 661,440 daily streams no one heard; giving every listen and every ticket a sealed proof of the real human — without violating privacy or the right to resell — is, on this front of the commerce of attention and leisure, **time2act**.

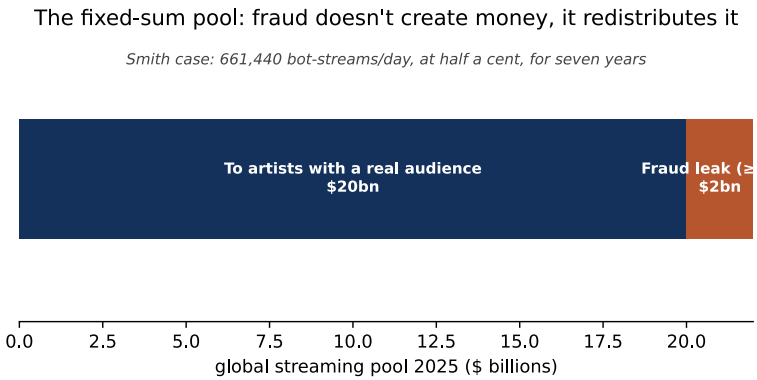


Figure 67.1 — Streaming music distributes a fixed-sum pool: in 2025, more than $22 billion (IFPI). Each fake stream invents no new money; it steals its share from artists with real audiences. The industry estimate puts the leak from fraud at some $2 billion a year — at least 10% of activity (Beatdapp). The Smith case: 661,440 bot streams a day, at half a cent, for seven years. With proof of life, listens without a human do not enter the distribution and the pool returns to real artists. Sources: IFPI Global Music Report 2026; Beatdapp; the New York prosecutors' indictment.

---

[1] U.S. Department of Justice / SDNY (via Music Business Worldwide) — *North Carolina man pleads guilty to AI-assisted music streaming fraud* (Michael Smith; $8.09 million forfeiture). https://www.musicbusinessworldwide.com/streaming-fraud-man-who-pocketed-8m-using-hundreds-of-thousands-of-ai-songs-streamed-billions-of-times-by-bots-pleads-guilty/ ↩

[2] Rolling Stone — *Inside the AI streaming-fraud case* (from 1,040 to 10,000 bot accounts; 661,440 plays/day; the emails). https://www.rollingstone.com/music/music-features/streaming-fraud-fake-streams-mike-smith-1235500686/ ↩

[3] Beatdapp / Music Business Worldwide — *Streaming fraud costs the global music industry ~$2bn a year* (≥10% of activity; an interested party's estimate). https://www.musicbusinessworldwide.com/streaming-fraud-costs-the-global-music-industry-2bn-a-year-according-to-beatdapp-now-its-partnering-with-beatport-to-combat-the-trend/ ↩

[4] Spotify / Music Business Worldwide — *The 1,000-stream royalty threshold* (announced Nov-2023; in force Apr-2024). https://www.musicbusinessworldwide.com/heres-exactly-spotifys-new-1000-annual-streams-royalty-policy-works123/ ↩

[5] The MLC / Billboard — *Mechanical Licensing Collective receives over $424 million in unmatched royalties* (the "black box"). https://www.billboard.com/pro/mechanical-licensing-collective-receives-over-424-million-unmatched-royalties-streaming-services/ ↩

[6] IFPI — *Global Music Report 2026* ($31.7 billion; streaming >$22 billion, 69.6%; "theft, plain and simple"). https://www.ifpi.org/global-music-report-2026-global-recorded-music-revenues-grow-6-4-as-record-companies-drive-innovation/ ↩

[7] IFPI (Victoria Oakley) — *Brazilian authorities disrupt JustAnotherPanel* ("fake streams… erode trust in our system"). https://www.ifpi.org/brazilian-authorities-launch-largest-ever-disruption-operation-against-major-streaming-fraud-platform-justanotherpanel/

[8] Federal Trade Commission — *FTC brings first-ever cases under the BOTS Act* (2021; >150,000 tickets; >$31 million in penalties). https://www.ftc.gov/news-events/news/press-releases/2021/01/ftc-brings-first-ever-cases-under-bots-act ↩

[9] Federal Trade Commission — *FTC sues Live Nation/Ticketmaster* (2025; 5 brokers, 6,345 accounts, 246,407 tickets; "turn a blind eye as a matter of policy"). https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-sues-live-nation-ticketmaster-engaging-illegal-ticket-resale-tactics-deceiving-artists-consumers ↩

[10] Wikipedia — *Taylor Swift–Ticketmaster controversy* (the Eras tour; 3.5 billion requests; ~12 million "entities" with bots). https://en.wikipedia.org/wiki/Taylor_Swift%E2%80%93Ticketmaster_controversy ↩

Telecommunications and Voice: Authenticating Who Really Calls

"Mr. Watson — come here — I want to see you."

— Alexander Graham Bell, laboratory notebook (10-3-1876)

In late January 2024, two days before the New Hampshire Democratic primary, thousands of phones rang. On picking up, an unmistakable voice — Joe Biden's — asked voters *not* to go vote and to "save their vote" for November. The caller ID showed a number known to belong to the party. It was all false: the voice was cloned with artificial intelligence and the caller ID, spoofed. The political operative who commissioned it, Steve Kramer, received a fine of **six million dollars**, and the U.S. Federal Communications Commission (**FCC**) issued the first rule declaring AI-voice *robocalls* illegal [2][3]. Nor was it the first time: back in 2019, the director of a British energy subsidiary had transferred some $243,000 after a call from his "boss," with the usual German accent — a synthetic voice, the first known cloned-voice fraud [10].

That New Hampshire afternoon contains this chapter's paradox: a voice of authority, **perfectly authenticated to the human ear, completely false on the wire**. Unlike chapter 16, which told of the visual impersonation of an executive on a video call, here the attack is on the **very infrastructure of telecommunications** — the SIM, the signaling network, the caller ID — and on the proof of **who really originates a voice**.

The number, the weak link

The most silent fraud does not touch the voice, but the number. In a **SIM swap**, the attacker gets the carrier to move your mobile number to a SIM he controls and, with that, intercepts the one-time codes that protect your accounts. In 2023 alone, the FBI's cybercrime office recorded more than a thousand complaints with losses of **$48.8 million** [1], and in the United Kingdom the anti-fraud service **Cifas** reported a **1,055%** increase in unauthorized SIM swaps in a single year [6]. The individual figures induce vertigo: T-Mobile was ordered to pay some **$33 million** — the award became public in 2025 — after a SIM swap enabled the theft of some $38 million in cryptocurrency from the entrepreneur Josh Jones, bypassing SMS verification [7]. As Cifas sums up, the mobile number has become **the weak link** of all digital identity based on SMS codes. And the preferred victims are the elderly: the U.S. trade commission found that 41% of older adults who lost ten thousand dollars or more to impersonation scams — as we have seen — said the first contact had been **a phone call**.

STIR/SHAKEN: a layer, not the solution

The industry has responded by authenticating the caller ID with the **STIR/SHAKEN** system, which cryptographically signs the originating number so carriers can block or label spoofed calls. But there is a huge hole: in April 2025 only **42.5%** of calls arrived signed, since the system **works only on IP networks** and any legacy segment breaks the signature [5]. As we have seen, American consumers received **52.8 billion robocalls** in 2024, some two hundred per adult [YouMail reference]. That is why the defense is shifting to the **network layer**: according to the **GSMA**, the global carriers' association, its Open Gateway initiative exposes signals such as "time since the last SIM change" so that a bank does not blindly trust an SMS [9]. But none of these layers answers the fundamental question: **who really originated this voice?**

The proof of existence: authenticating the human origin of the voice

This is where **t2p** contributes a complementary layer. It does not prevent someone from cloning a voice or taking over a SIM; **it seals the authenticity and origin of an act of identity at the moment it happens**. When a person records an instruction or authorizes a sensitive operation — authorizing a transfer, confirming identity to a bank, registering consent — the app computes on the device the **SHA-512 fingerprint** of the clip plus its metadata, anchors it to the **public blockchain** and attaches an **RFC 3161 timestamp** conforming to eIDAS and a **forensic GPS** fix, all in an **.asice** container.

The decisive piece is the **proof of life**, which attacks exactly the hole regulators point to: distinguishing a real human voice, captured live, from a synthesis. It is the same principle as the concept awarded by the trade commission itself in its *Voice Cloning Challenge*, **OriginStory**, which measures the biosignals of the throat and mouth to authenticate the **human origin of the voice at the point of creation**. Thus a legitimate authorization is sealed with liveness, time and place, so that a later cloned audio cannot pass itself off as the "original recording" or claim temporal precedence; the victim can prove they did **not** issue a certain instruction because no valid seal of theirs exists; and the timestamp fixes an incontrovertible order against a fabricated clip. STIR/SHAKEN authenticates the *number* on the network; t2p authenticates the *content and human origin* at the endpoint. They are layers that complete each other.

The limit must be said in all its rawness. "Garbage in, garbage on-chain": t2p seals the integrity and moment of a capture, but **it does not certify that its content is true or that the person is not being deceived at that instant**. It does not detect cloning inside the incoming call: if a scammer calls you with a cloned voice, t2p does not analyze it; it only works if the legitimate counterparty **proactively seals** their own act of identity — it is not a robocall filter or a substitute for STIR/SHAKEN. Nor does it prevent the SIM swap, which happens at the carrier; at most, it allows a system **not to depend** on an SMS as the sole proof. And if the attacker tricks the victim into recording and sealing an authorization herself, the chain will tell the truth about a lie: the seal will be valid and the act, fraudulent. Finally, like STIR/SHAKEN, it protects only those who use it at both ends, and it does not recover the money from a transfer already executed.

The promise and the peril

The promise is real on a terrain where the voice has stopped being proof. A layer that authenticates the human origin and the moment of an act of identity would give banks, companies and citizens a verifiable counterproof against the cloned voice, and would break the unhealthy dependence on SMS as the only key. Combined with STIR/SHAKEN on the network and the GSMA's APIs, it would complete a layered defense that today has holes everywhere.

The peril calls for the nuances we have stated. The first: do not confuse t2p with a call filter; it protects one's own act, but does not analyze someone else's. The second: it stops neither the SIM swap nor the coercion of the victim. And the third: adoption. Without use at both ends, there is no network effect. The proof of existence is a powerful complementary layer against voice impersonation, but it is not a silver bullet against all phone fraud.

It's time to act

For the **carriers**: complete STIR/SHAKEN, close the non-IP segments and apply the know-your-customer and know-your-upstream-provider principles the FCC imposed on Lingo.

For **banks and apps**: stop depending on SMS as the only proof and use network signals, such as the GSMA's SIM-change API.

For **regulators** (the FCC): prosecute illegal cloned voice, as in the New Hampshire case. And for **users**, especially the elderly: have a sealed proof attesting the human origin of their own voice, so that no clone can speak in their name.

A cloned Biden voice asked thousands of people not to vote, perfectly believable to the ear and utterly false on the wire. Giving every act of voice a sealed proof of its human origin — without forgetting that it stops neither the stolen SIM nor the deception — is, on this front where we can no longer believe what we hear, time2act.

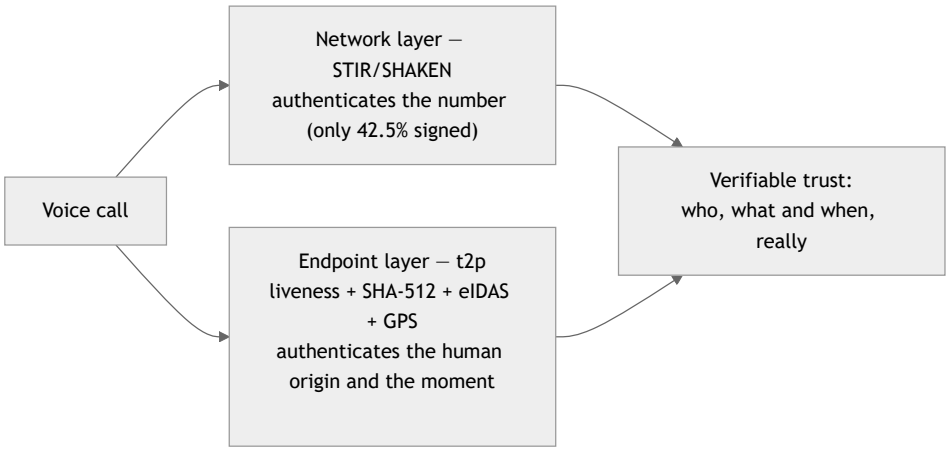


Figure 68.1 — The network layer (STIR/SHAKEN) authenticates the originating number — "from what number is the call coming?" — but only 42.5% of calls arrive signed (April 2025), because any non-IP segment breaks the signature. The endpoint layer (t2p: proof of life + SHA-512 + eIDAS seal + GPS) authenticates the human origin and the moment of the act — "who, what and when, really?". They are complementary. Context: 52.8 billion robocalls in the U.S. in 2024; $48.8 million stolen via SIM swap (FBI, 2023). Sources: TransNexus; FCC; FBI IC3.

---

[1] FBI Internet Crime Complaint Center (IC3) — *2023 Internet Crime Report* (SIM swapping: 1,075 complaints, $48.8 million in losses). https://www.ic3.gov/AnnualReport/Reports ↩

[2] NPR — *The AI deepfake Biden robocall and the FCC* (New Hampshire; Steve Kramer; $6 million fine). https://www.npr.org/2024/05/23/nx-s1-4977582/fcc-ai-deepfake-robocall-biden-new-hampshire-political-operative ↩

[3] Federal Communications Commission — *FCC makes AI-generated voices in robocalls illegal* (Declaratory Ruling, February 2024). https://www.fcc.gov/document/fcc-makes-ai-generated-voices-robocalls-illegal ↩

[4] Federal Communications Commission — *FCC settles Lingo Telecom robocall investigation* ($1 million; STIR/SHAKEN + know your customer and your upstream provider). https://www.fcc.gov/document/fcc-settles-lingo-telecom-robocall-investigation

[5] TransNexus — *SHAKEN statistics, April 2025* (only 42.5% of calls signed; the non-IP gap). https://transnexus.com/blog/2025/shaken-statistics-april/ ↩

[6] Cifas — *Fraudscape 2025* (+1,055% unauthorized SIM swaps in the United Kingdom). https://www.cifas.org.uk/newsroom/huge-surge-see-sim-swaps-hit-telco-and-mobile ↩

[7] SecurityWeek — *T-Mobile coughed up $33 million in SIM-swap lawsuit* (Josh Jones; ~$38 million in cryptocurrency). https://www.securityweek.com/t-mobile-coughed-up-33-million-in-sim-swap-lawsuit/ ↩

[8] Federal Trade Commission — *Winners of the Voice Cloning Challenge (OriginStory)* (authenticating the human origin of the voice at the point of creation). https://www.ftc.gov/news-events/news/press-releases/2024/04/ftc-announces-winners-voice-cloning-challenge

[9] GSMA — *Open Gateway / Mobile Identity APIs (SIM Swap API)* (the defense migrates to the network layer). https://www.gsma.com/solutions-and-impact/technologies/mobile-identity/mobile-identity-and-data-apis/ ↩

[10] The Next Web / The Wall Street Journal — *Fraudsters deepfake a CEO's voice* (2019; ~$243,000; the first known synthetic-voice fraud). http://thenextweb.com/news/fraudsters-deepfake-ceos-voice-to-trick-manager-into-transferring-243000 ↩

The Bank Fraud of the "Authorized" Payment

"Confidence is the indispensable basis of all sorts of business transactions. Without it, commerce [...] would, like a watch, run down and stop."

— Herman Melville, The Confidence-Man (1857)

The call comes from a number that looks like the bank's. The voice — professional, calm — warns of a "suspicious charge" and guides the victim, step by step, "to protect their savings": a transfer "to the safe account." Or the message from "the son": "Dad, my phone broke, send me a Bizum to this number." By the time the person realizes, the money is gone — and they sent it themselves, with their app, their codes and their fingerprint.

The legal knot is exactly this [2]. The European payments framework — PSD2, transposed in Spain by RDL 19/2018 — distinguishes two worlds. In the unauthorized transaction — someone else moves your money — the bank must refund the amount unless it proves fraud or gross negligence by the customer, and the burden of proving that the transaction was correctly authenticated is the bank's. But in the payment authorized under deception — what the British call APP fraud — the system's classic answer is devastating: "you authorized it." The perfect victim of modern fraud is not the one who loses their codes: it is the one who uses them convinced.

This chapter deals with the fraud that does not force the door: it convinces the victim to open it. The cloned SIM and the spoofed voice — the vehicles of the deception — are in chapter 68; the impersonated finance director ordering corporate transfers, in chapter 87; the relationship with the utility company, in chapter 33; institutional compliance, in chapter 16; the scam against the elderly, in chapter 85; off-platform payment in private sales, in chapter 31. The key question: when you made the payment yourself, deceived, what is left to argue about? The answer: diligence — what they showed you, what they warned you, what you saw — and diligence is dated facts: the certain date on both sides of the screen, with **t2p** — the triad of chapter 13: fingerprint, certain time, anchor [3]. The United Kingdom has just turned this battle of evidence into law.

The United Kingdom flips the board

Since October 7, 2024, the United Kingdom is the first country in the world with a mandatory reimbursement regime for authorized push payment fraud [1]. The regime — from the Payment Systems Regulator — obliges payment providers to reimburse victims: individuals, microenterprises and charities, on transfers over the Faster Payments and CHAPS systems, up to a limit of £85,000 — lowered from the initial 415,000 to align it with the British guarantee fund. Reimbursement within five business days — or up to thirty-five if investigation is needed. And the most ingenious piece: the cost is split 50% between the sending bank and the receiving bank — the bank that hosted the mule account pays too [1]. The regime replaces a voluntary code that had proven insufficient.

This book's reading: when reimbursement is the norm, the whole system becomes a dispute over diligence. Did the customer ignore clear warnings? Did the bank really show them — and when? Did the receiving bank open the mule account carelessly? Every answer is a dated fact on a screen or in a record. The British regime, without saying so, has created the need for proof of existence on both sides: whoever can prove what was shown, what was warned and what was done — and when — will have their share of the right; whoever cannot, will pay.

And in continental Europe, with the payments reform in progress [2], the question is not whether this model arrives, but when.

The proof of existence applied

The deception unfolds as a chain of moments — and every moment can be dated [3].

1. The deception documented

The fake technician's call, recorded — one's own conversation, lawful, as chapter 75 explains; the "safe account" text and the "son's" chat, captured with a certain date [3]. The deception stops being a story the

victim tells in shame: it is a dated sequence an analyst can retrace — and which, before the bank or the judge, establishes the context of the "authorized" payment: authorized, yes; free, no.

2. The session captured

What the victim was seeing when they paid: the app screen with the beneficiary and the warning — or the absence of a warning —, the website impersonating the bank, captured with the full image and the source code [3]. In a regime of shared diligence, this is the exact battlefield: the bank will say it warned; the customer, that they saw nothing. The screen of that moment, with a certain date, answers for both.

3. The dated claim

The clock runs from the first minute: the communication to the bank, captured on sending; the police report, scanned; every reply and every silence, with its date [3]. The deadlines of the British regime — reimbursement in 5 days, or 35 with investigation [1] — and those of any European claim are measured against these dates; and the whole .asice file, verifiable offline, reaches the complaints service, the supervisor or the court with the chronology closed.

4. The diligent bank's symmetry

The same mechanism protects the institution that does things right: the warnings displayed — the friction applied to the unusual payment, the new-payee warning — anchored in chained batches, with the same logic VeriFactu imposes on invoices [3]; and the record of the receiving account's onboarding, with its checks. When the regime allocates costs by diligence [1], the bank that can prove it warned — and the receiver that can prove it verified — have their defense; the one that can prove nothing, pays. The proof picks no side: it binds everyone.

The limits, stated clearly: the seal attests to screens, conversations and dates — it does not decide whether the negligence was gross or whether the warning was clear enough: that is for the law and the judge [2]. It does not recover money already moved — speed of reaction and cooperation between banks remain decisive. And prevention remains the first line: the bank never calls to ask for transfers, and no urgency is ever a good sign.

The promise and the peril

The promise: victims who do not have to choose between shame and defenselessness — the documented deception is claimable [1][3]; diligent banks that prove their warnings and do not pay for the sins of the careless; mule accounts more expensive to open, because the receiver answers too [1]; and a model — the British one — showing that APP fraud is not a fatality, but a problem of incentive allocation and proof.

Now for the perils.

The first is **the fraud of the fraud**. If reimbursement is automatic, someone will appear pretending to be a victim. The counterweight is the same file: the real victim has a dated, coherent sequence of deception; the pretender must fabricate it whole against a forensic GPS and seals that will not be backdated [3].

The second is **the bank's fortress**: the institution with perfect records against the customer with a memory. If only one side can prove, the regime tilts — that is why the customer's layer, their captures and their diary, is as important as the bank's, and the supervisor must weigh both.

The third is **the fear that paralyzes**. A regime too harsh on banks can end in infinite friction for everyone: blocked payments, universal suspicion. The balance is fine — clear, dated warnings, not obstacles; protection, not paternalism.

And the fourth is **reimbursement techno-solutionism**. Getting the money back does not undo the deception or catch the scammer. Pursuing the crime — and the mule networks — remains police work; the proof of existence helps it: every dated file is a map of the modus operandi.

It's time to act

Let those who trusted not have to prove their innocence empty-handed.

For you, the customer: if the deception comes, document while you react — the call, the text, the screen [3]; and remember the golden rule: the bank never calls to ask you for a transfer.

For sending banks: anchor your warnings and frictions [3] — in a regime of shared diligence [1], the warning you cannot prove does not exist.

For receiving banks: account onboarding with provable verification — when the mule account costs you money too [1], your front door is your first defense.

For European supervisors and legislators: the British model [1] already works; let the European payments reform [2] incorporate APP fraud reimbursement — and let the cost allocation be decided on dated evidence from both sides, not on the word of the biggest.

For police and prosecutors: each victim's .asice file [3] is a dated map of the network — aggregated, the maps catch mules and scammers.

Modern fraud does not steal codes: it steals trust. Let every deception at least leave its chronology — and let those who trusted not have to prove their innocence empty-handed. For everyone's savings, **time2act**.

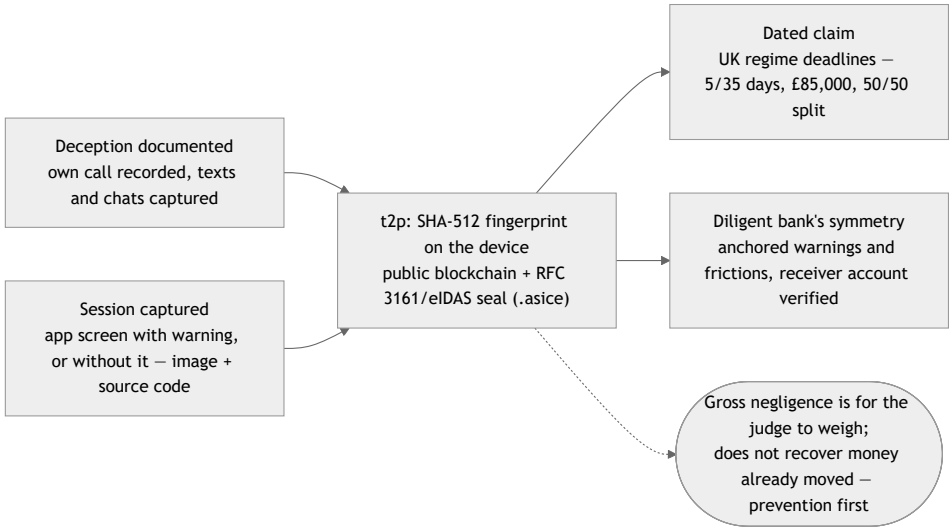


Figure 69.1 — The "authorized" payment fraud with proof of existence (sealed with t2p — chapter 13): the deception documented (own call recorded, texts and chats captured), the session captured (the screen with the warning — or without it), the dated claim (deadlines) and the diligent bank's symmetry (warnings and frictions anchored in batches; the receiving account's onboarding verified). Context: United Kingdom, the world's first mandatory reimbursement (7-10-2024; £85,000; 50/50 sender-receiver; 5/35 days); PSD2: in unauthorized transactions, the burden of proof is the bank's. Limits: gross negligence is for the judge to weigh; prevention first. Sources: PSR; PSD2/RDL 19/2018.

---

[1] Payment Systems Regulator (United Kingdom) — *APP fraud reimbursement protections* (mandatory regime since 7-10-2024, the world's first; individuals, microenterprises and charities; Faster Payments and CHAPS; £85,000 limit — lowered from 415,000 to align with the FSCS; 50/50 cost sharing between sending and receiving provider; reimbursement in 5 business days or up to 35 with investigation; replaces the voluntary CRM code); analyses by Freshfields, A&O Shearman and Thomson Reuters. https://www.psr.org.uk/information-for-consumers/app-fraud-reimbursement-protections/ · https://www.freshfields.com/en/our-thinking/briefings/2024/09/authorised-push-payment-fraud-a-new-mandatory-reimbursement-regime-for-uk-psps/ · https://www.aoshearman.com/en/insights/ao-shearman-on-fintech-and-digital-assets/the-uks-authorised-push-payment-app-fraud-reimbursement-scheme · https://www.thomsonreuters.com/en-us/posts/investigation-fraud-and-risk/app-fraud-uk/ ↩

[2] PSD2 (Directive (EU) 2015/2366, transposed in Spain by RDL 19/2018): in unauthorized payment transactions, refund by the provider unless there is fraud or gross negligence by the customer; burden of proof of correct authentication on the provider; strong customer authentication (SCA); the European payments reform (PSD3/PSR) in progress. ↩

[3] t2p — product manual (audio recording of one's own conversations; certified screenshot and web page capture — full image + source code; certify any file in batches; projects for an incident file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; the cloned voice and the SIM, in chapter 68; the impersonated CFO, in chapter 87; the personal black box, in chapter 75. ↩

Hardware: The Pedigree of a Chip That Could Not Fail

"Engineering design has as its first and foremost objective the obviation of failure."

— Henry Petroski, To Engineer Is Human (1985)

In a warehouse in Clearwater, Florida, between 2006 and 2010, a company called VisionTech imported integrated circuits from a single Chinese factory, made with "inferior and recycled" materials, and relabeled them with the brands of Intel, Texas Instruments and Motorola. Prosecutors found there a detail that says it all: **photographs of the owner's son applying car polish to dirty chips** so they would look new. The company claimed to do "internal testing," but it did not have a single quality-control expert. It sold **$15.8 million** in components labeled "suitable for military use," and its administrator ended up with thirty-eight months in prison — the first federal criminal conviction in the United States for trafficking in counterfeit integrated circuits [3].

Those recycled chips were not headed for toys. In 2010, the same kind of fake memory turned up in the mission computers of the **THAAD** anti-missile system; according to the defense agency, had they failed, **the missile itself would probably have failed** [1]. And a counterfeit memory chip was installed in more than four hundred cockpit display units of C-27J and C-130J transport aircraft, where a failure could leave the pilot's screen **blank in mid-flight** [4]. This chapter is not about the consumer-goods counterfeiting of chapter 29, but about a more technical and more dangerous lie: that of an electronic component's **pedigree** — its origin, its date, its chain of custody — which travels on forgeable labels and certificates into the heart of a weapons system.

The gray market and the pedigree lie

The scale of the problem is national. A year-long investigation by the U.S. **Senate Armed Services Committee** identified in 2012 some one thousand eight hundred occurrences of suspect counterfeit parts in the defense supply chain, involving **more than a million individual parts** and more than six hundred and fifty companies; more than **70% were traced to China** [1]. To show how easy it was, the congressional watchdog (**GAO**) created a fictitious company and bought sixteen military-grade components online: **all sixteen turned out to be fake** [2]. Most of these counterfeits — around **80%, according to industry databases** such as IHS and ERAI — are not copies built from scratch, but used chips ripped from discarded boards, polished and remarked. A single Shenzhen supplier, Hong Dark, introduced some **84,000 suspect parts** that reached military aircraft collision-avoidance systems and artillery shells [1]. The semiconductor industry association calculates that all this costs more than **$7.5 billion** a year [5], and the problem is not receding: the monitoring organization ERAI recorded in 2024 the highest number of suspect parts since 2015 [6].

Why detection is expensive and prevention at origin is cheap

What makes this fraud diabolical is that a recycled chip can be, on the outside, **identical** to an authentic one. As the CALCE research center at the University of Maryland explains, detecting it requires an expensive physical-chemical laboratory — X-rays, decapsulation, electrical tests — because the counterfeit lives inside. The aerospace standard **AS6171** standardizes these tests for parts that arrive **without a known chain of custody** back to the original manufacturer, and U.S. law holds the defense contractor responsible **no matter where the fake part entered**. But as the ERAI organization warns in its annual report, the fraud is atomized and hidden, and most cases are new occurrences. The sector's lesson is clear: **preventing at origin is far cheaper than detecting by testing**. The problem is that that origin travels in a PDF that can be rewritten.

The proof of existence: an unalterable pedigree

This is where **t2p** strikes at the heart of the fraud, which is not (only) a bad chip, but **a lie about the pedigree**. The idea is to seal every **custody event** — the marking by the original manufacturer, the incoming inspection, the AS6171 test, the acceptance report — computing on the device the **SHA-512 fingerprint** of the evidence artifacts: the test report with its X-ray images, the certificate of conformity,

the photo of the date code and lot number. This fingerprint is anchored to the **public blockchain** with an **RFC 3161 timestamp** conforming to eIDAS. The **proof of life**, at the moment the component's image is captured, prevents reusing an **old photo of an authentic part** to cover a fake lot, and the **forensic GPS** fixes **where** the inspection took place — the real accredited laboratory, not a "phantom lab." Everything travels in an **.asice** container that any actor in the chain can verify offline; if a date code or a certificate is altered, the fingerprint breaks. The chain of custody stops being a rewritable PDF and becomes an **unalterable, dated, geolocated pedigree record**.

The limit must be said in all its rawness, because here it is literal: "Garbage in, garbage on-chain." **t2p** seals the evidence and its moment, **not the physical truth of the silicon**. If an accredited laboratory correctly tests an authentic chip and then, downstream, someone **physically substitutes** that chip with a recycled one bearing the same marking, the chain will keep pointing to the legitimate seal: the attestation is honest, the part is not. The proof guarantees that *this report existed and has not been altered*, not that *the test was done well* or that *the part in my hand is the one that was tested*. That is why this front depends on a **reliable physical anchor** — DNA marking, a physically unclonable fingerprint in the chip itself, indelible serialization — that cryptographically binds the atom to the bit. Without this physical link, a bad actor can **certify lot A in good faith and deliver lot B**. The proof of existence raises the cost of the deception and leaves a forensic trail of who signed what and when — invaluable in litigation — but **it does not by itself detect a recycled chip**: that remains the job of physical testing, and a corrupt laboratory or a seal on the wrong part defeats the system.

The promise and the peril

The promise is one of national security. A sealed, dated, geolocated pedigree at every link would make it much harder to place a recycled chip in the heart of a missile or a cockpit display, because every false certificate would break the fingerprint and every responsibility would be tied to whoever signed it. Combined with physical testing and an anchor to the silicon, it would give the defense chain the traceability the law already demands but which today rests on forgeable papers.

The peril calls for the nuances we have stated. The first, do not confuse the seal with the physical test: the unalterable pedigree does not make a fake chip authentic, and without an anchor to the silicon the later substitution remains possible. The second, testing remains irreplaceable, and a corrupt laboratory defeats any chain. And the third, adoption across the whole chain, from the original manufacturer to the last subcontractor. The proof of existence is a powerful layer for the documentary pedigree, not a substitute for physical testing or for the anchor to the component.

It's time to act

For original manufacturers and authorized distributors: seal the marking and the date code at origin, so the pedigree is born unalterable. For prime contractors and subcontractors: demand the sealed pedigree as a condition, since the law holds you responsible wherever the fake part enters. For accredited laboratories: seal the test reports with proof of life and GPS, so no one reuses an old one. And for the Department of Defense and its agencies: bind the cryptographic pedigree to a physical anchor in the silicon, so the atom and the bit cannot be separated.

Chips polished with car wax ended up inside a missile that could not fail. Giving every component a sealed proof of its pedigree — never forgetting that physical testing and an anchor to the silicon are needed — is, on this front that touches everyone's security, time2act.

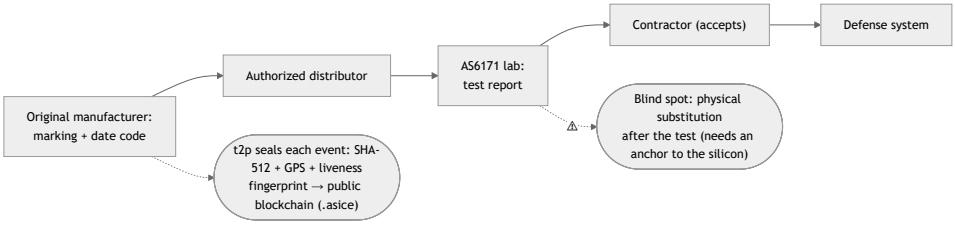


Figure 70.1 — Every custody event of a defense-grade component — marking by the original manufacturer, reception at the distributor, "origin" declared by the broker, AS6171 test report, contractor acceptance, installation — can be sealed with a SHA-512 fingerprint, an eIDAS timestamp,

forensic GPS and proof of life, anchored to the public blockchain. The blind spot (in red): the physical substitution of the part after *testing, which the seal does not detect without an anchor to the silicon (DNA, physically unclonable fingerprint). Context: ~80% of fake chips are recycled or remarked; the GAO bought 16 military parts online and all 16 were fake; a single broker introduced 84,000. Sources: U.S. Senate (2012); GAO; IHS/ERAI.*

---

[1] U.S. Senate Committee on Armed Services — *Inquiry into Counterfeit Electronic Parts in the Department of Defense Supply Chain* (2012; >1 million parts; >70% traced to China; THAAD; Hong Dark). https://www.armed-services.senate.gov/press-releases/senate-armed-services-committee-releases-report-on-counterfeit-electronic-parts ↩

[2] U.S. Government Accountability Office — *DOD Supply Chain: Suspect Counterfeit Electronic Parts on Internet Platforms* (GAO-12-375; 16 of 16 parts fake). https://www.gao.gov/products/gao-12-375 ↩

[3] Center for Public Integrity — *Counterfeit chips plague Pentagon weapons systems* (VisionTech; the polish on the chips; McCloskey, 38 months). https://publicintegrity.org/national-security/counterfeit-chips-plague-pentagon-weapons-systems/ ↩

[4] Air & Space Forces Magazine — *C-27J and C-130J parts linked to counterfeit mills* (L-3; blank screen; 27% defects). https://www.airandspaceforces.com/c-27jandc-130jpartslinkedtocounterfeitmills/ ↩

[5] Semiconductor Industry Association — *Winning the battle against counterfeit semiconductor products* (>$7.5 billion/year; ~11,000 jobs). https://www.semiconductors.org/resources/winning-the-battle-against-counterfeit-semiconductor-products/ ↩

[6] ERAI — *2024 Annual Report* (1,055 suspect parts; the highest since 2015). https://www.erai.com/erai_blog/3187/_2024_annual_report ↩

[7] U.S. Department of Justice (NDCA) — *Steve H.S. Kim sentenced for selling $3.5M in counterfeit and substandard parts* (42 months; nuclear submarine, missile system). https://www.justice.gov/usao-ndca/pr/california-resident-steve-hs-kim-sentenced-selling-35m-counterfeit-and-substandard

[8] SAE International — *AS6171: test methods for suspect/counterfeit EEE parts* (ISO/IEC 17025 accredited laboratories). https://www.sae.org/standards/content/as6171/7/

[9] CALCE (University of Maryland) — *Counterfeit electronics detection and avoidance* (a recycled chip can look identical to the authentic one). https://calce.umd.edu/counterfeit-electronics-detection-and-avoidance

[10] Defense Microelectronics Activity (DMEA) — *Trusted Foundry / Trusted IC* (trusted suppliers with a controlled pedigree; 2003, with the NSA). https://www.dmea.osd.mil/TrustedIC.aspx

Software: The Pedigree of the Code That Makes Everything Run

"The moral is obvious. You can't trust code that you did not totally create yourself."

— Ken Thompson, "Reflections on Trusting Trust," Turing Award lecture (1984)

Late March 2024. Andres Freund — a Microsoft engineer and PostgreSQL developer — is investigating a trifle almost anyone else would have let pass: the SSH logins on a test machine are taking about 500 milliseconds too long. Half a second. Pulling on that thread, Freund finds what no one wanted to find: a supremely sophisticated backdoor inside XZ Utils — a compression library present in almost all of Linux — designed to allow remote code execution through the SSH server. He reports it on March 28-29 to the Openwall security list; it receives the identifier CVE-2024-3094 and the maximum severity: 10.0 out of 10 [1].

Analysts agree on the retrospective diagnosis: had that half second of curiosity not triggered the investigation, the backdoor could have been an intrusion enabler bigger than the celebrated SolarWinds attack — potential access to any exposed Linux machine on the planet [1]. Banks, hospitals, power grids: everything that runs on Linux had been one step away from having a secret chamber opened inside it. The digital world was saved by 500 milliseconds and one stubborn engineer.

This chapter deals with the software supply chain: the code libraries that hold everything up, their versions and the volunteers who maintain them. The hardware pedigree — counterfeit chips — is in chapter 70; the records the attacker cannot erase, in chapter 57; the provenance of AI models, in chapter 79. The key question: how do you know the code you run — you, your bank, your hospital — is exactly what its authors published? The answer: with the code's pedigree — the public transparency logs like Sigstore/Rekor, which are already proof of existence in production [3], and the neutral anchoring layer of **t2p**: the SHA-512 fingerprint of every published version, with an RFC 3161/eIDAS timestamp and anchoring to the public blockchain, inside an .asice container [5]. The software industry has arrived on its own at this book's thesis; this chapter explains why — and what it still lacks.

The infiltrated bricklayer

When you open your banking app and see the encrypted-connection icon, you do not ask yourself who wrote the code that encrypts. The answer is: strangers. Modern software is a tower of pieces — "libraries" — that no one writes from scratch: every application takes blocks from open projects, often maintained by volunteers on a zero budget. XZ Utils was that: a humble compression piece, present in nearly every Linux distribution, tended for years by a single maintainer.

The attack was a masterpiece of social engineering, woven over several years [1]. A character calling himself "Jia Tan" appeared contributing useful fixes; fake accounts began pressuring the original maintainer with complaints and demands until, exhausted, he accepted help; and "Jia Tan" kept gaining permissions — first to modify code, then to publish releases.

Then he laid the poisoned brick, with a decisive subtlety: the malicious code was not visible in the public repository anyone can review. It hid in a build-process file and in obfuscated binary test files, inside the distribution packages — the "tarballs" — that distributions download to install the software [1]. Translation: the published version did not match the visible source code, and no standard repository review could detect it.

The evidentiary lesson is the chapter's turn: the attack exploited exactly a proof-of-existence gap — no one was verifying that what was published was what was visible. With the anchored fingerprint of every published package, checkable against the source code's fingerprint [5], the discrepancy would have been detectable by any machine in the world at the moment of download. No half-second hero would have been needed.

The worm with the keys to the kingdom

A year and a half later came the industrialization. On September 14, 2025, the npm ecosystem — the repository that feeds millions of web applications — suffered the first successful self-replicating worm in

its history: "Shai-Hulud," christened by the cybersecurity community after the sandworms of *Dune*. More than 500 compromised packages, including @ctrl/tinycolor — more than 2.2 million weekly downloads — and, in a perfect irony, packages from the npm account of CrowdStrike, one of the world's leading cybersecurity vendors [2].

The mechanics were as sophisticated as they were sinister: the worm harvested credentials from infected developers — Amazon, Google and Microsoft keys, located with secret-scanning tools — installed itself in their GitHub automations and, the lethal novelty, used each infected maintainer's publishing rights to publish itself inside their own packages: every victim became a propagator. A legitimate package on Monday could carry the worm on Tuesday, published with the same credentials. CISA — the U.S. cybersecurity agency — issued an alert on September 23; in November came a second wave [2].

The diagnosis, in this book's key: if xz was the infiltrated bricklayer who took years, Shai-Hulud is the theft of publishing identity at scale. The problem is no longer only who the maintainer is — identity, which chapter 91 covers — but what gets published with their keys and when. The software supply chain's question is, literally, this book's question: what existed, exactly, and at what moment?

The answer that is already proof of existence

The good news deserves to be spelled out: the industry has begun to respond with the very architecture this book defends. Sigstore — the open software-signing infrastructure — has become the default method for ecosystems like Kubernetes and registries like npm, Maven and PyPI; and its heart is Rekor: a public, append-only transparency log — nothing in it can be deleted or rewritten — which already holds more than 101 million signature entries from more than 33,000 open source projects, from Kubernetes to CPython [3]. Every signed release leaves a public, dated record, and npm and PyPI already accept provenance attestations signed this way [3].

The most elegant defensive use is log monitoring: a maintainer can watch Rekor and receive an alert if a publishing signature for their package appears there that they did not make — detecting key theft by the public trail it leaves, as Trail of Bits and the OpenSSF demonstrated in December 2025 [3]. Against a worm like Shai-Hulud [2], it is exactly the missing alarm: the fraudulent release is recorded where everyone can see it — including the rightful owner, within minutes.

And here the thesis moment: an append-only transparency log with dated entries is, by definition, proof of existence. The software community, pushed by its attackers, has arrived on its own at the solution this book proposes for the rest of society. What t2p adds is the neutral, transversal anchoring layer: SHA-512 fingerprints of every published version — tarballs included, the lesson of xz [1] — with an RFC 3161/eIDAS timestamp and anchoring to a public chain [5], so that the code's pedigree depends on no single infrastructure and is verifiable with the same tool as the pedigree of a document, a photo or a chip (chapter 70).

The regulatory context closes the circle: the European Cyber Resilience Regulation — Regulation (EU) 2024/2847, in force since December 2024 — imposes on manufacturers security across the digital product's whole life cycle, with vulnerability notification obligations applying from September 11, 2026 and full application on December 11, 2027 [4]. Software traceability stops being good practice and becomes law.

The proof of existence applied

Any update, any server — and, before installing it, a one-second check against a log no one can rewrite.

1. The version with a pedigree

Every published version of a library — the full package, the tarball — gets its SHA-512 fingerprint anchored with a certain date [5], checkable against the fingerprint of the visible source code. At install time, the package manager compares the fingerprint of what it downloaded with the anchored one: if they do not match — the discrepancy that hid the xz backdoor [1] — it does not install, and the alarm rings on its own. The code's pedigree, like the chip's in chapter 70: not a vendor's promise, but a chain verifiable by anyone before trusting it.

2. The watched release

The maintainer anchors their publishing chain — which versions they have published, when, with which keys — and monitors the public logs [3]: any release with their credentials that they did not make triggers the alarm in minutes, not weeks. Against the worm that publishes itself with stolen keys [2], the infection window shrinks from days to the instant of detection — and the whole community can see, in the log, which versions are legitimate and which are not. The watchman stops being one company: it is everyone.

3. The sealed inventory

Every digital product carries — and the European framework will progressively demand it [4] — its inventory of components: which libraries, which versions (the so-called SBOM, the software's list of ingredients). An inventory sealed per version, with a certain date [5], turns incident response into a query: when the next xz or the next Shai-Hulud erupts [1][2], every organization can know in minutes — and prove to the regulator [4] — whether or not it runs the poisoned piece, and since when. Guesswork becomes verification.

4. The forensic timeline

After an attack, the question is worth millions: when did the malicious code get in, which versions carry it, who published it? With anchored publishing chains and inventories [3][5], the forensic investigation works on a non-repudiable timeline — the same logic as the records of chapter 57, applied to code. And responsibility is apportioned justly: the manufacturer that deployed the clean version can prove it; the one that ignored the dated alert is equally on record.

The limits, said plainly: the seal attests that a package existed and has not been altered — it does not attest that the code is good or safe. The backdoor the legitimate author publishes with their own key is perfectly anchored — xz is half an example: "Jia Tan" was a maintainer [1]. Code review, auditing and key security remain the first line; and the transparency log detects the fraudulent release, it does not prevent it — the fast response remains human.

The promise and the peril

The promise: a software supply chain where every piece carries a verifiable pedigree [3][5], discrepancies like xz's sing on their own [1], worms like Shai-Hulud meet alarms within minutes [2] and incident response is a query instead of an archaeology [4]. The world that runs on code, running on attested code.

The perils, with the same clarity.

The first is **the seal's false security**. Anchored does not mean harmless. Malicious code published by its legitimate author with their legitimate keys gets impeccably anchored [1]; the pedigree attests to origin and integrity, not goodness. Whoever sells the seal as a security audit will be creating the worst vulnerability: unjustified trust.

The second is **the burden on volunteers**. The world's digital infrastructure rests on unpaid maintainers — the single, exhausted maintainer of xz is the pattern, not the exception [1]. If the pedigree becomes one more bureaucratic obligation on their backs, they will walk away — and the problem will get worse. The tools must be automatic and free for open source, and funding the digital commons is part of the security solution [4].

The third is **the centralization of the log**. If the whole world's software pedigree depends on a single transparency log, we have created a single point of failure and of power [3]. The answer is the whole book's: plurality of logs, cross-anchoring to neutral public chains [5] and no single gatekeeper of the code's truth.

And the fourth is **chain techno-solutionism**. No fingerprint fixes the political economy of the problem: patient state attackers [1], ecosystems that reward publishing fast, infinite dependencies no one reviews. The pedigree makes the fraud visible; reducing it also takes review, funding and culture. Cryptography is the scaffolding, not the building.

It's time to act

If everything runs on code, let the code have a pedigree.

For open source maintainers: sign and anchor every release — tarball included [1][3][5] — and monitor the log: let no release with your keys catch you by surprise [2][3].

For companies that deploy software: demand a verifiable pedigree for every piece, keep the inventory sealed per version [4][5] and check fingerprints before installing — Andres Freund's half second [1], automated.

For registries and ecosystems (npm, PyPI and company): provenance attestations by default [3], fingerprint verification on download and public alarms for anomalous releases [2].

For European regulators: the Cyber Resilience Regulation [4] arrives at just the right moment — let its application recognize the transparency log and the anchored fingerprint as means of compliance.

For governments and big tech: fund the digital commons — the exhausted maintainer is vulnerability number one [1]; no cryptography replaces a decent salary for those who hold up everyone's infrastructure.

The world was saved one March dawn by half a second of latency and an engineer's stubbornness [1]. Let the next backdoor not depend on luck: let it run into a pedigree that exposes it on the first download. For the code that holds everything up, **time2act**.

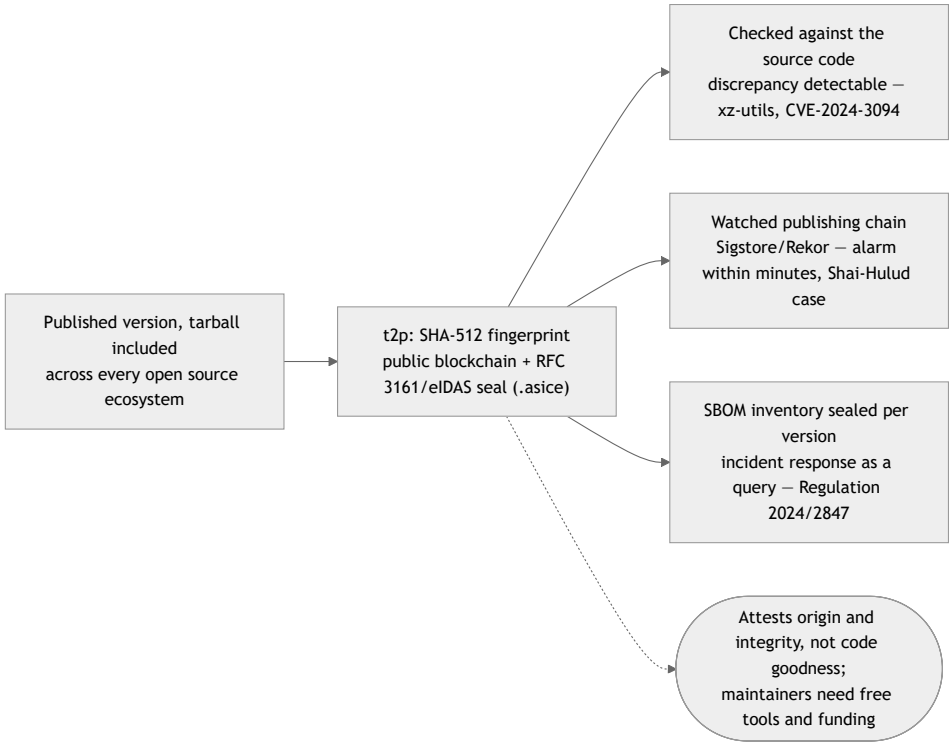


Figure 71.1 — The software supply chain with proof of existence: every published version is anchored (SHA-512 fingerprint of the tarball, RFC 3161/eIDAS timestamp, public blockchain, .asice) and checked against the source code (the xz-utils discrepancy, detectable); the publishing chain watched in the transparency logs (Sigstore/Rekor: >101 million entries — proof of existence in production) triggers the alarm on releases with stolen keys (the Shai-Hulud worm); the inventory (SBOM) sealed per version makes incident response a query (Regulation (EU) 2024/2847). Limits: the seal attests to origin and integrity, not the code's goodness; maintainers need free tools and funding. Sources: CVE-2024-3094; CISA (Shai-Hulud); Sigstore/Rekor; CRA.

---

[1] Wikipedia — *XZ Utils backdoor* (CVE-2024-3094; CVSS 10.0; Andres Freund, Microsoft engineer and PostgreSQL developer, investigating ~500 ms of extra SSH latency; report to Openwall on March 28-29, 2024; "Jia Tan's" social engineering campaign with fake accounts pressuring the maintainer; code hidden in build-to-host.m4 and binary test files inside the tarballs — invisible to repository review; potential greater than SolarWinds; caught before mass deployment); Akamai — *XZ Utils Backdoor — Everything You Need to Know*; JFrog — *XZ Backdoor Attack CVE-2024-3094*; Datadog Security Labs — *The XZ Utils backdoor*. https://en.wikipedia.org/wiki/XZ_Utils_backdoor · https://www.akamai.com/blog/security-research/critical-linux-backdoor-xz-utils-discovered-what-to-know · https://jfrog.com/blog/xz-backdoor-attack-cve-2024-3094-all-you-need-to-know/ · https://securitylabs.datadoghq.com/articles/xz-backdoor-cve-2024-3094/ ↩

[2] Palo Alto Unit 42 — *"Shai-Hulud" Worm Compromises npm Ecosystem in Supply Chain Attack* (September 14, 2025; the first self-replicating worm on npm; >500 packages; @ctrl/tinycolor with >2.2 million weekly downloads; CrowdStrike packages; AWS/GCP/Azure credential harvesting with TruffleHog; persistence via GitHub Actions; propagation with the maintainers' publishing rights); CISA — *Widespread Supply Chain Compromise Impacting npm Ecosystem* (alert of September 23, 2025);

StepSecurity — *Shai-Hulud: Self-Replicating Worm Compromises 500+ NPM Packages*; Datadog — *The Shai-Hulud 2.0 npm worm* (second wave, November). https://unit42.paloaltonetworks.com/npm-supply-chain-attack/ · https://www.cisa.gov/news-events/alerts/2025/09/23/widespread-supply-chain-compromise-impacting-npm-ecosystem · https://www.stepsecurity.io/blog/ctrl-tinycolor-and-40-npm-packages-compromised · https://securitylabs.datadoghq.com/articles/shai-hulud-2.0-npm-worm/ ↩

[3] Sigstore — *Rekor: Software Supply Chain Transparency Log* (public append-only log; >101 million entries from >33,000 projects: Kubernetes, CPython, Istio, urllib3; default signing method for Kubernetes, npm, Maven and PyPI; provenance attestations on npm and PyPI); OpenSSF / Trail of Bits — *Catching malicious package releases using a transparency log* (December 2025; monitoring Rekor to detect fraudulent releases with stolen credentials). https://docs.sigstore.dev/logging/overview/ · https://www.sigstore.dev/ · https://openssf.org/blog/2025/12/19/catching-malicious-package-releases-using-a-transparency-log/ · https://blog.trailofbits.com/2025/12/12/catching-malicious-package-releases-using-a-transparency-log/ ↩

[4] OJEU / BOE — *Regulation (EU) 2024/2847, of October 23, 2024* (Cyber Resilience Act: horizontal cybersecurity requirements for products with digital elements; in force in December 2024; notification obligations from September 11, 2026; full application on December 11, 2027; security across the whole life cycle); Atalaya GRC — *CRA: fechas y obligaciones*; Red Seguridad — *Cyber Resilience Act: una nueva era para los fabricantes*. https://www.boe.es/buscar/doc.php?id=DOUE-L-2024-81720 · https://atalayagrc.com/normativas/cra/ · https://www.redseguridad.com/especialidades-tic/normativa-y-certificacion/cyber-resilience-act-cra-una-nueva-era-para-los-fabricantes_20260204.html ↩

[5] t2p — product manual (SHA-512 fingerprint; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; version fingerprints and batch manifests; neutral, transversal anchoring, complementary to sector transparency logs). See chapters 13 and 15; the hardware pedigree, in chapter 70; the unerasable records, in chapter 57. ↩

Fishing and Minerals: The Origin That Sank with the Ship

"For all at last return to the sea — to Oceanus, the ocean river, like the ever-flowing stream of time, the beginning and the end."

— Rachel Carson, The Sea Around Us (1951), final sentence

On December 17, 2014, two ships of the Sea Shepherd organization began chasing a poaching trawler called the *Thunder*, which had spent years fishing toothfish illegally in Antarctic waters. The pursuit lasted **one hundred and ten days and ten thousand nautical miles**, all the way to the economic zone of São Tomé and Príncipe. And there, instead of surrendering, the crew did something revealing: **they deliberately sank their own ship**, with the hatches lashed open, to destroy the evidence of the illegal catch. Three officers were convicted of forgery and environmental crimes, with a fine of fifteen million euros [Thunder case reference]. The proof of origin, literally, sank with the ship.

That ship sinking itself to erase its trail contains this chapter's question. In chapter 29 we saw the traceability of agricultural products on land against deforestation; here the pattern is the same — sealing the origin — but in two new and sometimes bloody domains: the **sea** and the **conflict-mineral mines**. The question is no longer "does this coffee come from deforested land?", but "was this fish caught in a closed zone, and by whom?" and "did this tantalum finance a militia?"

One in five fish, and 90% "clean" mineral

The scale of marine fraud is immense. Illegal, unreported and unregulated fishing moves between **eleven and twenty-six million tonnes** a year — close to one in five wild-caught fish — worth between $10,000 and $23,500 million [1]. And it is nearly invisible: a **Global Fishing Watch** study published in *Nature* in 2024 showed, crossing satellite imagery with artificial intelligence, that **three out of four of the world's industrial fishing vessels emit no public position signal at all** [3]. When the fish reaches the plate, moreover, it is often not what it claims: the organization **Oceana** found that a third of U.S. seafood was mislabeled, and up to **87%** of what was sold as "red snapper" was another species [2]. And behind all of it lies the horror: the Environmental Justice Foundation's investigation documented Burmese fishermen **locked in cages** on the Indonesian island of Benjina, slaves of a fleet whose fish ended up in Western supermarkets [7].

The mineral parallel is just as dark. The **3TG conflict minerals** — tin, tantalum, tungsten and gold — finance armed groups in the Democratic Republic of the Congo, regulated by the U.S. Dodd-Frank Act and the European regulation of 2017. Diamonds have their own system, the **Kimberley Process**, created in 2003, which reduced the trade in conflict diamonds from close to 15% in the nineties to less than 1% [6]. But the problem is present, not historical: the armed group M23 took the Rubaya coltan mine in 2024, netting some **$800,000 a month** [10].

When traceability becomes a laundromat

And here comes this chapter's most important warning, because it strikes at the heart of the solution. A traceability system is not infallible; it can become a **laundromat**. The organization **Global Witness**, in its 2022 report *The ITSCI Laundromat*, showed that in two mining zones of South Kivu between **80% and 90%** of the minerals labeled "clean" by the zone's only traceability provider actually came from illicit sources [5]. The trick was simple and devastating: the "clean" tag was applied to mineral **already mixed**. A good-provenance label on bad-provenance mineral is worse than no label, because it confers false confidence. And the risk of centralizing traceability in a single company is real too: Everledger, the pioneer in tracking diamonds with blockchain, went into insolvency in 2023 [9]. The lesson is clear: **the problem is not having a certificate, but where and when it is born**.

The proof of existence: sealing the origin at the point of extraction

This is where **t2p** attacks the fraud at its exact origin. The idea is to seal the catch or the lot **at the moment and place of extraction**, before it can be mixed or relabeled. On board, with the net freshly

hauled, or at the mouth of the artisanal mine, the operator takes a photo of the catch with the **proof of life** of the fisherman or the miner; the device locally computes the **SHA-512 fingerprint** of the image and the metadata — declared species, weight, fishing gear, vessel or mine identifier — embeds the **forensic GPS** and an **RFC 3161 timestamp** conforming to eIDAS, and anchors the fingerprint to the **public blockchain**, generating an **.asice** container that survives without coverage on the high seas or in a mine gallery.

The effect is direct against the three frauds. Against **backdating**: no one can claim that a fish caught in a closed zone was caught in a legal one, because the GPS-and-time seal is prior and unalterable. Against **species swapping**: the sealed photo binds a visual identity to the lot at origin. And against **lot laundering** — the laundromat case: every sack of coltan is sealed with its mine's geolocation **before** any mixing, so that blending militia mineral breaks the chain of fingerprints instead of inheriting a clean label. The European catch certificate or the Kimberley certificate stops being a PDF validated late and gains a cryptographic anchor of origin.

The limit must be said in all its rawness, because here it is the most dangerous of all. "Garbage in, garbage on-chain," in its extreme form: **laundering at the point of sealing**. t2p does not by itself solve what Global Witness uncovered. It does not verify the species or the legality of the zone: the fingerprint guarantees that *this photo, at these coordinates, at this time, taken by this living person* existed — not that the fish is the declared species (that takes DNA) or that the coordinates lie outside a closure (that takes crossing them with maps of protected areas). It does not prevent **the sealer's corruption**: if the one sealing is the operator of the poaching vessel or the buyer at the militia mine, they can seal illegal catches with true metadata and false declarations. It does not cover **the dark vessel**: the 75% of the fleet that emits no signal would simply seal nothing, and against those who choose not to participate, satellite surveillance is still needed. And it does not fix the photo's physical attribution to the object — a photo can show one lot while another is shipped — so it contributes more combined with physical marking, such as nanotracers or DNA, anchored to the same fingerprint. Its real value is **closing the window of backdating and documentary swapping**, not guaranteeing by itself that the declared origin is true.

The promise and the peril

The promise touches the environment and human rights at once. Sealing the origin at the point of extraction, before any mixing, would make it far harder to slip fish from a closed season or coltan from a militia into a global supply chain, and would give the catch and Kimberley certificates an anchor they lack today. In a world where one in five fish is illegal and where a "clean" label can hide blood mineral, sealing the origin is protecting the sea and people at the same time.

The peril calls for the nuances we have stated, and here they are especially grave. The first, the laundromat: a seal on already-mixed mineral confers false confidence, which is why the sealing must happen **at origin and before mixing**, never after. The second, the sealer's corruption, which no cryptography solves without independent auditing. And the third, the dark vessel, against which the satellite remains necessary. The proof of existence is a powerful layer for anchoring the origin, not a substitute for DNA, for geospatial verification, or for auditing whoever seals.

It's time to act

For **flag states and the European Union**: demand the catch certificate with an origin anchor sealed at the point of fishing, not validated late in an office. For **mineral importers**: seal every lot at the mouth of the mine, before any mixing, to prevent the laundromat. For the **Kimberley Process and the brands** (electronics, jewelry): recognize the seal of origin as primary evidence. And for the **independent auditors and satellite surveillance**: keep watching the dark vessel and the honesty of whoever seals, because no label is worth anything without them.

A poaching ship sank itself to erase the origin of its catch. Giving every catch and every lot of mineral a proof sealed at the point of extraction — never forgetting that DNA, the satellite and the audit are needed — is, on this front that touches the sea and human dignity, time2act.

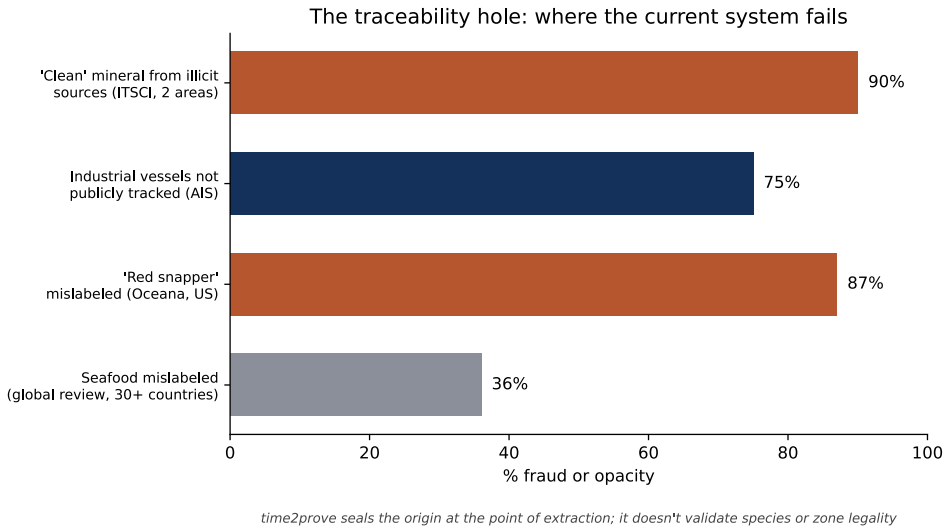


Figure 72.1 — Fraud and opacity by vertical: three out of four industrial fishing vessels emit no public position signal (75%, Global Fishing Watch); a third of U.S. seafood is mislabeled — up to 87% of "red snapper" (Oceana); a global review across more than thirty countries finds 36% mislabeling; and up to 90% of "labeled clean" mineral in two zones of South Kivu came from illicit sources (Global Witness). t2p seals the origin at the point of extraction, before mixing or relabeling; what the seal does NOT validate is the species or the legality of the zone. Sources: Global Fishing Watch; Oceana; Global Witness.

---

[1] Agnew et al. (PLOS ONE) — *Estimating the worldwide extent of illegal fishing* (11-26 million tonnes; $10,000-$23,500 million/year). https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0004570 ↩

[2] Oceana — *Oceana study reveals seafood fraud nationwide* (33% mislabeled; "red snapper" 87%). https://oceana.org/reports/oceana-study-reveals-seafood-fraud-nationwide/ ↩

[3] Global Fishing Watch / Nature — *Satellite mapping reveals extensive industrial activity at sea* (~75% of vessels with no public tracking). https://globalfishingwatch.org/press-release/new-research-harnesses-ai-and-satellite-imagery-to-reveal-the-expanding-footprint-of-human-activity-at-sea/ ↩

[4] European Commission — *Illegal fishing (IUU): catch certificates and CATCH* (Reg. (EC) 1005/2008; CATCH mandatory since 10-Jan-2026). https://oceans-and-fisheries.ec.europa.eu/fisheries/rules/illegal-fishing_en

[5] Global Witness — *The ITSCI Laundromat* (up to 80-90% of "clean" mineral from illicit sources). https://www.globalwitness.org/en/campaigns/natural-resource-governance/itsci-laundromat/ ↩

[6] Kimberley Process / World Diamond Council — *The Kimberley Process fact sheet* (from ~15% in the 90s to <1%; ~99% of production). https://www.worlddiamondcouncil.org/wp-content/uploads/2022/02/The-Kimberley-Process-Fact-Sheet.pdf ↩

[7] Environmental Justice Foundation — *Thailand's Seafood Slaves* (slavery on board; the island of Benjina). https://ejfoundation.org/reports/thailands-seafood-slaves ↩

[8] Science Advances (Park et al.) — *Illuminating Dark Fishing Fleets in North Korea* (>900 vessels; >164,000 t; >$440 million). https://www.science.org/doi/10.1126/sciadv.abb1197

[9] Ledger Insights — *Everledger bankruptcy* (the risk of centralizing diamond traceability in a single company). https://www.ledgerinsights.com/everledger-bankruptcy-esg-blockchain-traceability/ ↩

[10] The Africa Report — *Rubaya coltan mine at the heart of M23 financing* (~$800,000/month for an armed group). https://www.theafricareport.com/375904/drc-rwanda-rubaya-coltan-mine-at-the-heart-of-m23-financing/ ↩

Philanthropy: Where Does My Euro Go

"Anyone can give money; but to give it to the right person, to the right extent, at the right time, with the right motive, and in the right way — that is not for everyone, nor is it easy."

— Aristotle, Nicomachean Ethics, II, 9 (1109a)

After the 2010 Haiti earthquake, the American Red Cross raised almost **five hundred million dollars**, while promising entire neighborhoods of new houses to a devastated people. In 2015, the journalists Justin Elliott, of ProPublica, and Laura Sullivan, of public radio NPR, walked the ground to see the result. They found **six permanent houses** completed in the whole country. On paper, the organization claimed to have "provided a home to more than 130,000 people," but that figure included people who had attended a repair seminar or received a temporary shelter already in ruins [1]. One detail summed it all up: a senior Haitian engineer earned $42,000 while an expatriate project manager earned $140,000.

Those six houses contain this chapter's question, which is the donor's, not the beneficiary's. In chapter 49 we saw how to prove that aid **reaches** those who need it; here the question is prior and complementary: **where does my euro go between the click and the impact, and is the impact they sell me real?** Because between the donation and the cause there is a dark stretch — fundraisers, expenses, subcontractors, inflated figures — where far more is lost than is stolen in the final distribution.

Where does my euro go?

The opacity of this stretch is systemic and, at times, criminal. As we have seen, in 2015 the U.S. trade commission and all fifty states sued four fake cancer charities, controlled by the Reynolds family, for swindling **more than $187 million**: less than **3%** reached patients, and the telemarketers kept 85% or more of every donation, which was spent on cruises, jet skis and concert tickets [2]. But you do not need outright fraud to get lost along the way: according to the New York attorney general's annual report, professional fundraisers kept **25% of every dollar** in 2024, and in half the campaigns the charity received **less than half** of what was raised [4]. And the problem is growing: the FBI's cybercrime office put fraud from fake charities and disaster campaigns at some **$96 million** in 2024, with artificial intelligence manufacturing the appearance of legitimacy [3].

Faced with this, the sector has learned a counterintuitive lesson. In 2013 the heads of the three big NGO evaluators — **GuideStar, Charity Navigator and the BBB Wise Giving Alliance** — signed an open letter, *The Overhead Myth*, warning that looking only at the percentage of administrative expenses is a **bad indicator**: it pushes NGOs into "starvation," into not investing in the systems that would make them effective [6]. The right question, as the effectiveness evaluator **GiveWell** argues, is not how much an NGO spends on itself, but **how many lives it changes per dollar** — and it calculates that the best ones can avert a death for between three thousand and five thousand five hundred dollars [7]. For the donor, as Daniel Borochoff, founder of **CharityWatch**, sums up, the guide should not be marketing but **auditable financial transparency**.

When the donation launders a reputation

There is a fraud subtler than the money kind: that of **reputation**. In December 2021, New York's Metropolitan Museum removed the Sackler family name from seven galleries. The Sacklers, owners of Purdue Pharma — the manufacturer that aggressively promoted OxyContin in the middle of the opioid epidemic — had been "among our most generous donors." As the journalist **Patrick Radden Keefe** documents in his book *Empire of Pain*, that philanthropy worked as **reputation laundering**: million-dollar donations to museums that bought a facade of virtue while the family never directed funds to treating the addiction its company had spread. The donation, here, does not steal money from the donor; it steals the truth from the public.

The proof of existence: sealing the promise and the evidence of impact

This is where **t2p** offers a layer that takes the claims out of the organization's word and anchors them in verifiable evidence. It does not certify that an NGO is honest, but it seals **evidence at every node of the money's journey**. First, the **promise**: at the instant of donating, the amount, the restricted purpose and

the published promise are summarized in a **SHA-512 fingerprint**, anchored to the **public blockchain** and sealed with the eIDAS timestamp. This fixes unalterably *what was promised* and *what* the fund was restricted to, preventing the promise from being rewritten later — the Haiti pattern, where a "new community" became a "seminar." Second, every **disbursement** — a paid invoice, a transfer to a subcontractor — is sealed with its amount and date, creating a chained trail of actual overhead versus promised. And third, the **evidence of impact** — the photo of the works, the signed certificate, the deliverables count — is captured with **forensic GPS**, a **timestamp** and **proof of life**, and packaged in an **.asice** container verifiable offline.

The effect for the donor is that they can **audit** whether the 75% said to reach the cause matches the sealed disbursements, and that an impact photo carries coordinates, time and proof of live capture — not an image recycled from another disaster. An independent evaluator like CharityWatch, a regulator or the donor themselves can verify the container without depending on the NGO's server.

The limit must be said in all its rawness, because here it is subtle. "Garbage in, garbage on-chain": the seal guarantees that a datum existed and when, not that the money was well spent. The **last-mile problem** is structural: when the NGO pays a local supplier in cash, that transaction leaves the verifiable trail. An impact certificate can be **sealed, geolocated and lying at the same time**; six real cases can be photographed and presented as proof of "130,000 helped," if the declared metric is rigged — the proof of existence does not judge the *definition* of the claim, only its existence. It does not detect reputation laundering or the dirty origin of the donor's money: sealing a Sackler donation makes it more auditable, not cleaner. And it does not replace the judgment of effectiveness — whether it is three thousand dollars per life or a useless project — which demands GiveWell-style evaluation and human auditing. **t2p** provides **integrity of the file**, not **merit** of the program.

The promise and the peril

The promise is to restore the trust that sustains all philanthropy. A layer that seals the restricted promise, every disbursement and the evidence of impact would give the donor the ability to audit their euro, would make it much harder to inflate beneficiaries or recycle photos, and would reward honest NGOs with provable transparency instead of penalizing them for a misunderstood overhead. In a sector that lives on trust, making the file unalterable protects both the donor and the organizations that do their work well.

The peril calls for the nuances we have mentioned. The first, the last mile: the trail ends where someone stops capturing honest evidence. The second, do not confuse integrity with truth: a rigged metric, sealed, is still rigged. The third, the seal neither launders dirty money nor judges effectiveness. And the fourth, selective capture: if the NGO chooses what it seals, the trail will have holes. The proof of existence is a powerful tool for the journey's transparency, not a certificate of honesty or a substitute for impact evaluation.

It's time to act

For donors: demand the restricted promise and the evidence of impact sealed, and focus on cost per outcome, not on overhead. For NGOs: seal every disbursement and every proof of impact, turning transparency into a provable advantage. For evaluators and regulators (CharityWatch, the attorneys general): recognize the sealed container as the basis for independent audit. And for fundraising platforms (crowdfunding): bind every campaign to a verifiable promise and evidence, so that a fake page cannot collect in a victim's name.

Five hundred million dollars produced six houses, and no one knew until two journalists told the story. Giving every donation a sealed proof of its journey — never forgetting that impact evaluation and human audit are still needed — is, on this front that sustains generosity, time2act.

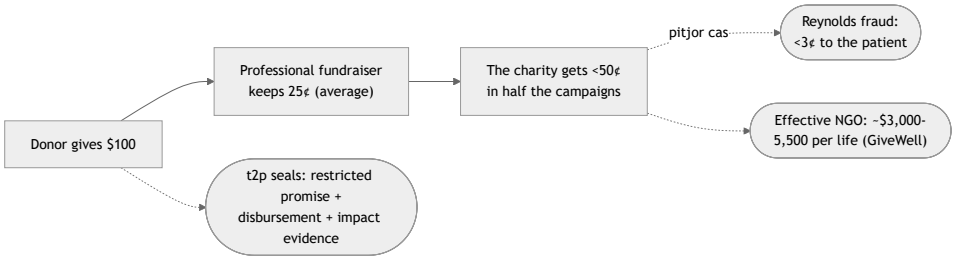


Figure 73.1 — The journey of a donated dollar, with its real leaks: the professional fundraiser keeps 25 cents on average (New York attorney general, 2024), and in half the campaigns the charity receives less than half; in the Reynolds fraud, less than 3 cents reached the patient; in the Haiti project, 40% was lost to overhead. As a counterpoint, an effective NGO averts a death for about $3,000-5,500 (GiveWell). t2p seals three nodes — the restricted promise, every disbursement and the evidence of impact (GPS + time + proof of life) — turning "trust the report" into auditable evidence. Limit: it does not judge whether the declared metric is honest or the program's effectiveness. Sources: New York attorney general; FTC; ProPublica; GiveWell.

---

[1] ProPublica / NPR (Justin Elliott and Laura Sullivan) — *How the Red Cross raised half a billion dollars for Haiti and built six homes* (2015). https://www.propublica.org/article/how-the-red-cross-raised-half-a-billion-dollars-for-haiti-and-built-6-homes ↩

[2] Federal Trade Commission — *FTC, all 50 states and D.C. charge four cancer charities with bilking over $187 million* (the Reynolds family; <3% to patients). https://www.ftc.gov/news-events/news/press-releases/2015/05/ftc-all-50-states-dc-charge-four-cancer-charities-bilking-over-187-million-consumers ↩

[3] FBI Internet Crime Complaint Center — *PSA: charity and disaster-relief fraud* (2024; ~$96 million; AI boosts apparent legitimacy). https://www.ic3.gov/PSA/2025/PSA250116 ↩

[4] New York Attorney General (Letitia James) — *Pennies for Charity 2025* (fundraisers keep 25%; <50% in half of campaigns). https://ag.ny.gov/press-release/2025/attorney-general-james-releases-annual-pennies-charity-report ↩

[5] CBS News — *CBS News investigates Wounded Warrior Project spending* (from $1.7 to $26 million on conferences; executives dismissed in March 2016). https://www.cbsnews.com/news/cbs-news-investigates-wounded-warrior-project-spending/

[6] GuideStar, Charity Navigator and BBB Wise Giving Alliance — *The Overhead Myth* (2013; overhead as a bad indicator). https://learn.guidestar.org/news/news-releases/2013/2013-06-17-overhead-myth ↩

[7] GiveWell — *How much does it cost to save a life?* (between ~$3,000 and $5,500 per life). https://www.givewell.org/how-much-does-it-cost-to-save-a-life ↩

[8] CharityWatch (Daniel Borochoff) — *Our charity rating process* (≥75% to program; ≤$25 to raise $100). https://www.charitywatch.org/our-charity-rating-process

[9] CBS News / Patrick Radden Keefe (*Empire of Pain*) — *The Met removes the Sackler name* (philanthropy as "reputation laundering"). https://www.cbsnews.com/news/metropolitan-museum-of-art-sackler-family-name-removed-purdue-pharma/

[10] Federal Trade Commission — *Reported fraud losses hit $12.5 billion in 2024* (fraud context). https://www.ftc.gov/news-events/news/press-releases/2025/03/new-ftc-data-show-big-jump-reported-losses-fraud-125-billion-2024

Black Boxes: The Data the Manufacturer Said Did Not Exist

"Facts are stubborn things; and whatever may be our wishes, our inclinations, or the dictates of our passions, they cannot alter the state of facts and evidence."

— John Adams, defense in the Boston Massacre trials (1770)

On April 25, 2019, in Key Largo, Florida, a Tesla with Autopilot engaged caused a fatal crash. When the victim's lawyers demanded the car's data, Tesla first replied that it **did not exist**; then it told the police the data was "corrupted," and invoked an **auto-delete function that did not exist**. It was all false: its own servers had received a **collision snapshot** — video, data from the car's bus and from the recorder — about **three minutes after the impact**. An outside expert, Alan Moore, recovered everything by making a **bit-by-bit image of the memory** of the control unit. "Tesla's engineers said it couldn't be done... and people outside Tesla did it." In August 2025, a Miami jury ordered Tesla to pay **$243 million**, of which 200 in punitive damages [1][2].

That tug-of-war over a datum contains this chapter's question, distinct from those of chapter 20 (evidence of atrocities for a tribunal) and chapter 57 (the records of a computer intrusion). Here the datum is **machine telemetry** — speed, braking, object classification, flight parameters — and the typical conflict is not "we were hacked," but **"the car says X, the driver says Y"** and, above all, **who controls and trims the manufacturer's data**. Because the problem, as we shall see, is not that data is missing: it is that whoever holds it can make it disappear.

The black box exists; the problem is who controls it

Modern recorders are prodigious. An aircraft black box survives an impact of 3,400 g, 1,100 degrees of fire for half an hour and water pressure at six thousand meters, and records more than eighty-eight parameters [5]. A car's recorder now captures the twenty seconds before a crash ten times per second [7]. The data is there. The problem is **custody and selection**. In the fatal running-over of Elaine Herzberg by an autonomous Uber in Tempe in 2018, the telemetry revealed that the system **detected the pedestrian 5.6 seconds before** but never managed to classify her — object, vehicle, bicycle — while emergency braking was disabled; the black box captured not a breakdown but an **algorithmic indecision** millisecond by millisecond [3]. And the Cruise case in San Francisco showed the danger in its purest form: after a robotaxi dragged a pedestrian six meters, the company showed the California traffic authority a **trimmed video that ended before the dragging**; the **DMV suspended its permits** for misrepresenting safety information [4]. When the one who controls the data is the same party under investigation, **the choice of what gets shown becomes the central problem**.

When the box does not reach the end, or does not settle the cause

But the black box has two limits no technology can erase. The first is that it only helps if it records to the end. In the Jeju Air crash in South Korea in December 2024, with 179 dead, both the voice recorder and the data recorder stopped about four minutes before impact, leaving the decisive stretch outside the box [9].

The second limit is deeper still: the box fixes the physical facts, but does not settle the interpretation. In the 1999 EgyptAir 990 crash, with the same authentic, intact data, the American authority concluded a deliberate action by the copilot, while the Egyptian one attributed it to a mechanical failure [10]. The bits were the same; their meaning, in dispute. The blockchain is no expert witness.

The proof of existence: sealing the black box

This is where **t2p** contributes a layer of **evidentiary integrity**, not just of capture. It seals the black box so that no one can later say "this is not the original data." At the moment the recorder closes an event — a collision snapshot, a block of flight data or a drone's landing log — the device computes the **SHA-512 fingerprint** and anchors it to the **public blockchain** with an **RFC 3161 timestamp** conforming to

eIDAS. Only the fingerprint is published; the raw data remains private, but its existence and exact content at that instant are proven.

This directly counters the cases we have seen: if the fingerprint of the complete snapshot is anchored three minutes after the impact, a later "different" datum will not match the seal, and a truncated video like Cruise's will show a fingerprint incompatible with the full clip's. The **forensic GPS** and the timestamp for drones and vehicles make it possible to refute the claim "I was not in a restricted zone," as in the Cape Canaveral case, where the drone's own logs contradicted the operator [8]. The **proof of life** at the moment of extraction certifies that an identified person pulled the data on such a date, sealing the chain of custody from the first minute.

Three limitations must be noted. First, "garbage in, garbage on-chain": the seal does not correct the source datum; if a system misclassifies a pedestrian, it only certifies that the error existed at that moment. Second, **it does not settle interpretation**: as the EgyptAir 990 case shows, the proof would have armored the bits but the expert dispute over what they mean would remain intact. Third, **it does not record what the machine never captured**: in Jeju Air the boxes stopped recording four decisive minutes, and the seal only closes off the gap. Finally, the anchoring **must happen on the device and immediately**; if the manufacturer controls the moment of sealing — as it controls the snapshot today — it could seal an already modified version. The value is real only with **automatic, on-device, instantaneous sealing**, outside the discretionary control of the party under investigation.

The promise and the peril

The promise is to close the asymmetry that cost $243 million to prove against Tesla; a black box sealed at origin would make it impossible to deny, corrupt or trim the data after the fact, and would give investigators and courts an unassailable chain of custody from minute zero, protecting both the victims and the honest manufacturers, who could prove their data is the original. In a world where crashes are decided by algorithms, making the black box unalterable is restoring trust in what really happened.

The peril calls for the nuances we have stated. First, do not confuse integrity with truth or with cause: the seal armors the bits, but does not settle their meaning. Second, it does not fill the gaps of what the sensor did not capture. And third, most important, if the sealing is not immediate and on the device, the manufacturer can still curate the data before sealing it. The proof of existence is a powerful layer of evidentiary integrity, not an expert witness or a substitute for a complete recording.

It's time to act

For manufacturers (of cars, aircraft, drones): seal every data block on the device and at the instant it closes, outside their own discretionary control. For regulators (transport safety boards, aviation, traffic): require sealing at origin as part of recorder homologation. For investigators and courts: seal the extraction with proof of life, so the chain of custody is unassailable. And for operators: data that protects them whether they are right or not, so that the truth of the crash does not depend on who controls the server.

A manufacturer said a datum did not exist, and an outside expert found it on its servers. Giving every black box a proof sealed at the moment of birth — never forgetting that the seal neither settles the cause nor fills in what was not recorded — is, on this front where machines decide, time2act.



Figure 74.1 — The recorders (the aircraft black box with more than 88 parameters and now up to 25 h of voice; the car's recorder with the 20 s before the crash; a drone's flight log) capture the data. t2p does not improve the capture: at the instant the recorder closes an event, it computes its SHA-512 fingerprint on the device and anchors it to the public blockchain with an eIDAS seal, so that denying the data (Tesla), trimming it (Cruise) or altering it leaves a fingerprint that does not add up. Limit: the seal does not correct an erroneous datum, does not settle the expert dispute over the cause (EgyptAir 990), and does not fill in what the sensor did not record (Jeju Air, where the recording cut out 4 minutes early). Sources: NTSB; California DMV; Benavides v. Tesla.

[1] NPR — *Tesla ordered to pay $243 million in fatal Autopilot crash* (Benavides v. Tesla; $200 million in punitive damages; 33% liability). https://www.npr.org/2025/08/02/nx-s1-5490930/tesla-autopilot-crash-jury-240-million-florida ↩

[2] Electrek — *Tesla withheld data and misdirected police in Autopilot crash* (the collision snapshot at +3 min; the NAND image; the expert Alan Moore). https://electrek.co/2025/08/04/tesla-withheld-data-lied-misdirected-police-plaintiffs-avoid-blame-autopilot-crash/ ↩

[3] National Transportation Safety Board — *Collision Between an Uber ATG Test Vehicle and a Pedestrian (HAR-19-03)* (5.6 s; emergency braking disabled). https://www.ntsb.gov/investigations/Pages/HWY18MH010.aspx ↩

[4] California DMV — *Statement on Cruise LLC suspension* (24-Oct-2023; the dragging footage, omitted). https://www.dmv.ca.gov/portal/news-and-media/dmv-statement-on-cruise-llc-suspension/ ↩

[5] National Transportation Safety Board — *Cockpit Voice Recorders and Flight Data Recorders* (survivability; ≥88 parameters). https://www.ntsb.gov/news/Pages/cvr_fdr.aspx ↩

[6] Federal Register / FAA — *25-hour Cockpit Voice Recorder requirement for new aircraft* (2026; from 2 to 25 hours). https://www.federalregister.gov/documents/2026/02/02/2026-02110/25-hour-cockpit-voice-recorder-cvr-requirement-new-aircraft-production

[7] eCFR — *49 CFR Part 563: Event Data Recorders* (pre-crash capture from 5 s to 20 s at 10 Hz). https://www.ecfr.gov/current/title-49/subtitle-B/chapter-V/part-563 ↩

[8] DroneXL — *Drone pilot sentenced over Cape Canaveral Space Force base flights* (the geofencing logs contradict the operator). https://dronexl.co/2025/08/28/drone-pilot-sentenced-cape-canaveral-space-force-base/ ↩

[9] CNN — *Jeju Air black boxes stopped recording before the crash* (the 4 decisive minutes, outside the box). https://www.cnn.com/2025/01/11/asia/south-korean-airline-black-boxes-stopped-recording-intl-hnk/index.html ↩

[10] National Transportation Safety Board — *EgyptAir 990 Aircraft Accident Brief (AAB-02/01)* (the same data, opposite causes). https://www.ntsb.gov/investigations/AccidentReports/Reports/AAB0201.pdf ↩

The Personal Black Box: Preventive Evidence for Victims and for the Innocent

"An ounce of prevention is worth a pound of cure."

— Benjamin Franklin, Pennsylvania Gazette (1735)

On March 20, 2024, nothing visible to the naked eye happened. No siren, no headline. That day, the Cometa system — the state service that records every movement of abusers wearing anti-abuse ankle bracelets — was changing hands: from Telefónica to the joint venture formed by Vodafone and Securitas. In the migration, the positioning histories prior to that date became inaccessible to the courts [1]. Over the following months, in courtrooms across Spain, prosecutors trying to prove breaches of restraining orders ran into the void: the annual report of the State Prosecutor General's Office, made public in September 2025, acknowledges "numerous acquittals and provisional dismissals" due to the loss of that data, in a system with some 4,500 active bracelets, and speaks of a potential unprotection of the victims [1][2]. The government had known of incidents since January 2024 [3]. The black box that was supposed to protect thousands of women failed in silence, and part of the evidence died with it.

This chapter deals with the personal black box: the preventive certification of your own life — the audio of a threatening conversation, the dashcam clip, the dated photo — sealed on the device in case one day you are a victim or you are falsely accused. The manufacturers' black boxes are in chapter 74; the judicial chain of custody, in chapter 77; the insurance claim, in chapter 28; documenting online harm, in chapters 83 and 84. The key question: how can an ordinary person keep the truth safe without becoming anyone's spy? The answer is **t2p**: the content never leaves the phone; only the SHA-512 fingerprint, with the RFC 3161/eIDAS timestamp and the anchor to the public blockchain inside an .asice container, is fixed forever [18]; and revealing it or not is always your decision. The neutrality is radical: the same tool serves the victim who documents — the everyday case — and the innocent who defends themselves — the rare but devastating one. The proof has no side.

When the state's black box fails

The Cometa case is worth a close look, because it is an entire parable. The country's most sensitive evidentiary infrastructure — the one that records abusers' movements to protect victims of gender violence — depended on a database migration between contractors. In the change from Telefónica to the Vodafone-Securitas joint venture, the positions prior to March 20, 2024 became inaccessible to the judicial authorities [1]. When a prosecutor needed to show that an abuser had approached his victim before that date, the datum was no longer there: hence the acquittals and provisional dismissals the Prosecutor's Office ended up acknowledging in its annual report [2]. The government had had notice of incidents since January 2024 [3]; the handling cost the minister a parliamentary reprimand and a public apology to the victims, without ruling out penalties for the contractor [4]. The subsequent legal analysis summed it up bluntly: a system "in crisis" [5].

The lesson is not partisan, it is structural: when the evidence lives in a single centralized repository, a change of contract, a migration error or an administrative decision can erase it. No malice is needed; fragility is enough. The proof of existence inverts the design: if every position and every alert had been anchored at origin — SHA-512 fingerprint and timestamp, at the very instant of being produced [18] — the content could still have been lost, but it would have been impossible to dispute what existed, what was recorded and when. No defense could have been built on doubt about whether the data ever existed.

And here is born the question that gives this whole chapter its meaning: if this can happen to the state, with 4,500 bracelets and a multi-million contract, what can happen to the citizen who, facing a conflict, has only their word?

The right to record your own life

The base rule was set by the Constitutional Court four decades ago, and it was born, precisely, from a recording submitted in a labor dispute. STC 114/1984, of November 29, holds that recording a conversation you take part in does not violate the secrecy of communications of Article 18.3 of the Constitution — "there is no secret for the one to whom the communication is addressed" — and that one's

own recording is admissible as documentary evidence, with one limit: not invading the other person's intimate sphere [6]. The Supreme Court has confirmed it consistently [7]. Your phone can be, legally, your witness.

The red line is just as sharp. Recording other people's conversations — those you are not part of — with technical devices of listening or recording is the crime of discovery and disclosure of secrets: Article 197.1 of the Criminal Code punishes it with one to four years' prison and a fine, and if the material is disseminated, the penalty rises to two to five years [8]. The personal black box is not a hidden microphone in someone else's home: it is the record of your own life, of the conversations you hold and of what happens in front of your windshield.

The windshield subcase deserves its own map, because the European mosaic is instructive. In Spain, the Spanish Data Protection Agency accepts the dashcam with conditions: recording limited to the incident — about twenty seconds before and after — preferably without sound and minimizing third parties' data; within these limits, the recording counts as evidence, covered by legitimate interest and effective judicial protection [13]. Germany admits it as evidence limited to the moment of the crash; France makes the driver responsible for the data processing; and Portugal applies the harshest policy in Europe: Law 58/2019 bans dashcams to the point that you cannot even carry one in the car [14]. The right to your own evidence exists, then, but it is narrow and demands scrupulous respect for others' privacy. It is exactly t2p's design: only fingerprints on the chain, content on the device, voluntary disclosure before the judge [18].

The recording is no longer believed

Two forces have collided this decade: it has never been so easy to record everything, and never so easy to fake everything. With current technology, convincing voices are cloned from three to five minutes of clean audio, and some systems make do with thirty seconds [9]: any fragment of conversation can be imitated by a machine that reproduces the speaker's timbre, rhythm and imperfections.

The courts know it, and they have moved. In proceedings, challenging the authenticity of an audio shifts the burden of proof onto whoever submits it; and the Supreme Court's recent case law no longer accepts the generic challenge — the "it could be fake" — but demands a technical, substantiated challenge, with admission criteria growing stricter: forensic analysis, a complete chain of custody, attested integrity, cross-checking against other evidence [10].

For the citizen, the consequence is uncomfortable: your truthful recording, on its own, is worth less every day. The day after a conflict, the other side can always sow doubt — "it's a pile of bits, edited who knows when" — and the expert report that resolves it is expensive, slow, and every generation of synthesis makes it harder [9][10].

The structural way out is chapter 15's: not detecting the falsehood after the fact, but attesting authenticity beforehand. A recording with its SHA-512 fingerprint sealed at the moment of capture and anchored to the public blockchain [18] arrives at the courthouse with the key question already answered: this file, exactly this one, existed on that date and has not been touched since. Reasonable doubt switches sides: whoever wants to deny it will have to explain how a cryptographic seal is forged backward in time.

The black box you already carry

The personal black box is not surveillance: it is your own certified memory, activated when you choose. A button before a tense meeting; a car mode that keeps only the incident; a photo that is sealed as it is taken. The content stays encrypted on the device; only the fingerprint travels [18]. It is not about recording your entire life, but about opening a controlled channel for those fragments that can decide your fate.

The state, in fact, is already pointing this way. The official app AlertCops, of the State Secretariat for Security, includes an SOS button that, pressed five times in under six seconds, sends the security forces your location and ten seconds of audio; victims of gender violence register with the AVDG code for automatic notice to the competent forces [15]. It is the institutional recognition that an immediate, dated audio can break the circle of word against word.

The victim of harassment — at work, from neighbors, at home — finds the first use here: accumulating over months her own sealed conversations, covered by STC 114/1984 [6]. When she decides to report, she

will not carry scattered memories: she will carry a verifiable chronology.

The falsely accused innocent finds the second. It is a statistically rare case — gender-violence reports ending in a final conviction for false reporting are 0.0082% of almost three million reports, according to the Prosecutor's Office series [11][12], and this book underlines it so as to feed no myth — but for whoever suffers it, it is devastating. A sealed diary — where you were, what you said — can be the only possible defense against a defamation or a mistaken accusation.

And the driver finds the third, against the staged-crash rackets: insurers detect some 400,000 fraud attempts a year, with a potential cost of €2 billion; six in every hundred declared injuries hide fraud, and staged crashes are 3.2% of those declared but 18% of the cost of investigated claims [16]. In July 2025, a single ring ended with 32 arrested over seven staged crashes in eight months [17]. The dashcam sealed at origin, within the AEPD's limits [13], is the private citizen's black box on wheels.

The proof of existence applied

An ordinary person, an ordinary phone, and a decision: keeping the truth safe in case it ever has to be shown.

1. The victim who documents

A woman suffering repeated threats records the conversations she takes part in, covered by STC 114/1984 [6]. Every night, her phone computes the SHA-512 fingerprint of the files, adds the RFC 3161/eIDAS timestamp and anchors it to the public blockchain [18]. The content does not leave the phone: no clouds, no third parties, no leak risk. The day she reports, her lawyer submits a chronology of evidence with a certain date that no generic challenge can knock down [10]. If the abuser deletes his messages or denies the calls, her record still exists — and it depends on no contractor and no database migration [1]. It complements, does not replace, the official channels: AlertCops and 016 [15].

2. The innocent who defends himself

A man accused of having said or done something at a specific moment produces his sealed diary: the audio of the real meeting, the dated location, the original message. It is not a tool of suspicion against anyone — final convictions for false reporting are minuscule [11][12] — but low-cost insurance against the worst day of a life. The key is technical: the seal was created before the conflict, and no one fabricates a retroactive alibi against a blockchain [18].

3. The black box on wheels

A driver with the camera configured the way the AEPD recommends — only the incident, about twenty seconds before and after, no sound [13] — runs into a suspicious sudden stop: two coordinated cars, four occupants, injuries declared the next day. The clip sealed at origin — fingerprint and time anchored in the same minute [18] — dismantles the ring's script [16][17]: against a cryptographic seal, no coordinated story holds. For the insurer, a faster expert assessment; for the honest driver, the end of "you against four witnesses."

4. The handshake contract

The freelancer who closes a job over the phone, the tenant who agrees on a repair with the landlord, the family that arranges the care shifts for an elderly father: their own conversations, legally recorded [6] and sealed on the spot [18]. If all goes well, no one will ever listen to them. If someone goes back on their word, the certain date turns "I never said that" into an untenable position.

It must also be said, with the same clarity, what the seal does not do: it attests to when a file existed and that it has not been altered, but not that what is said in it is true; it does not make lawful a recording that was not — recording other people's conversations remains a crime under Article 197 [8]; and it does not relieve the judge of weighing the evidence in context. The personal black box is armored memory, not a verdict.

The promise and the peril

The promise is threefold: that no victim ever again loses a trial because the evidence lived on someone else's server and vanished with a contract [1][2]; that no innocent depends on their word alone; and that the truthful recording recovers the evidentiary value that voice cloning and synthesis are eroding [9][10]. But this is, perhaps, the chapter where the perils demand the most frankness in the whole book.

The first is **the society of suspicion**. If everyone records everyone all the time "just in case," intimate trust dies: a society of mutual black boxes is a society surveilled from within. That is why the personal black box must be defensive and exceptional: activated with intent, not by default; recording your life, not other people's; and the fingerprints-only design [18] guarantees that nothing can be disseminated that has not been formally disclosed. The tool is for the courtroom, not for the networks.

The second is **the criminal line and the intimate sphere**. The law protects one's own recording (STC 114/1984 [6]) and punishes the interception of other people's conversations (Article 197 of the Criminal Code [8]). No seal launders an unlawful capture, and dissemination outside the proceedings can be a crime even if the recording was lawful. The border between legal evidence and public exposure must be engraved in the tool's very design.

The third is **revocability and oblivion**. The fingerprints of your life must be invalidatable at your request, and there can be no central registry of personal fingerprints. Immutability protects the evidence; it cannot become a life sentence upon one's own memory.

And the fourth is **the excuse not to fix the state**. The personal black box does not exempt the administration from making Cometa, AlertCops or 016 work [1][15]. That the citizen has to armor their own evidence is a symptom, not an ideal: the promise is just only if it coexists with institutions that also do their job.

It's time to act

The proof of your life cannot depend on memory or on a maintenance contract.

For you, citizen: know your right — recording conversations you take part in has been legal since STC 114/1984 [6]; recording other people's, a crime [8]. Activate the black box with intent — the difficult meeting, the car journey, the incident — and seal on the device [18]. And never disseminate: disclose only before the judge.

For victims of gender violence and harassment: register in AlertCops with the AVDG code [15]; document with your own sealed recordings and seek legal advice. A verifiable chronology is worth more than a thousand painful memories.

For lawyers, judges and experts: demand and weigh the capture seal; an anchored fingerprint resolves from the outset the authenticity challenge that today makes proceedings longer and costlier [10].

For administrations: anchor public evidentiary systems at origin — let no migration ever again erase the protection of 4,500 women [1][2][3] — and grant qualified evidentiary value to sealed citizen evidence.

For insurers: reward the AEPD-compliant dashcam with a seal at origin [13]: less fraud [16][17], shorter assessments and better treatment for honest customers.

For legislators: establish the presumption of authenticity for evidence sealed at capture; guarantee its revocability; and not one step toward mandatoriness — the personal black box is a right, never a duty.

The state's black box failed one March 20 with no one watching [1]. Yours depends on no contract, no migration and no company: only on you and on mathematics. For each person's truth, **time2act**.

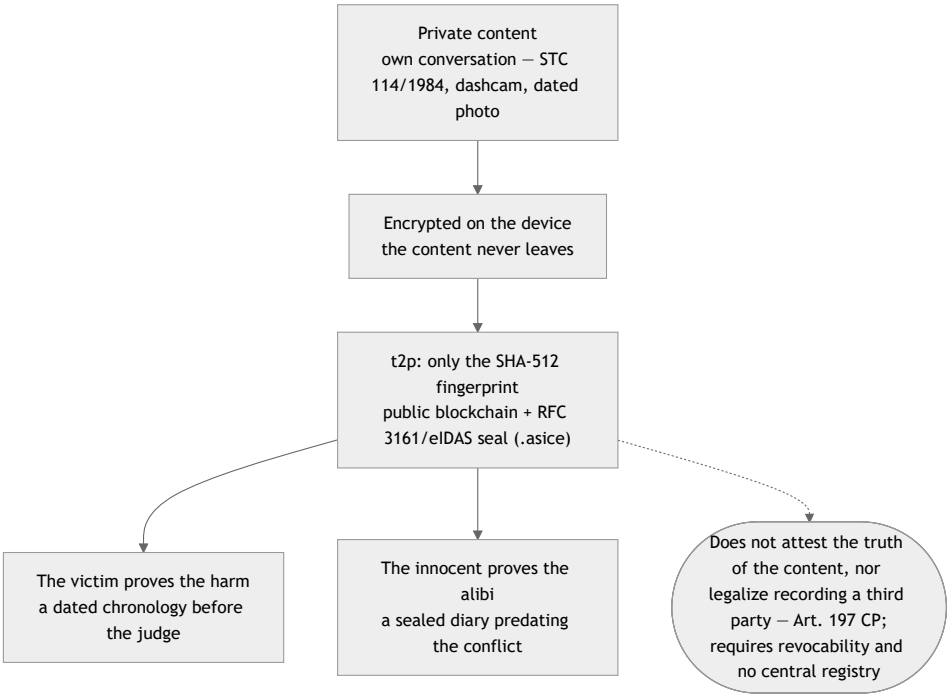


Figure 75.1 — The personal black box: the audio of one's own conversation (STC 114/1984), the dashcam clip (AEPD limits) or the dated photo are encrypted and NEVER leave the device; only the SHA-512 fingerprint with the RFC 3161/eIDAS timestamp is anchored to the public blockchain (.asice). Disclosure is the holder's decision, before the judge. Two symmetrical outcomes: the victim who proves the harm and the innocent who proves the alibi. Limits: it does not attest the truth of the content, does not legalize recording a third party (Art. 197 CP), requires revocability and no central registry. Sources: STC 114/1984; State Prosecutor General's Office (Cometa; false reports 0.0082%); AEPD; UNESPA.

---

[1] Euronews — *Un fallo informático en las pulseras antimaltrato en España absuelve a decenas de agresores* (September 18, 2025; Telefónica → Vodafone+Securitas joint venture migration; data prior to March 20, 2024 inaccessible; State Prosecutor General's annual report; ≈4,500 active bracelets). https://es.euronews.com/2025/09/18/un-fallo-informatico-en-las-pulseras-antimaltrato-en-espana-absuelve-a-decenas-de-agresore ↩

[2] Diario Socialista — *Fallo en pulseras «antimaltrato» provoca absoluciones y pérdida de pruebas en casos de violencia machista* (September 18, 2025; "numerous acquittals and provisional dismissals"). https://diariosocialista.net/2025/09/18/fallo-en-pulseras-antimaltrato-provoca-absoluciones-y-perdida-de-pruebas-en-casos-de-violencia-machista/ ↩

[3] OKDiario — *El Gobierno reconoce ahora que conocía desde «enero de 2024» los fallos en las pulseras antimaltrato.* https://okdiario.com/espana/pulseras-gobierno-reconoce-ahora-que-sabia-desde-enero-2024-fallos-pulseras-antimaltrato-15967835 ↩

[4] La Nación (agencies) — *Redondo aborda en el Senado la gestión de pulseras para maltratadores tras ser reprobada en el Congreso* (September 2025; apologies to the victims; possible penalties for the contractor). https://www.lanacion.com.ar/agencias/redondo-aborda-este-miercoles-en-el-senado-la-gestion-de-pulseras-para-maltratadores-tras-ser-nid28092025/ ↩

[5] Hay Derecho — Nuria Barona Martínez, *Pulseras antimaltrato: análisis jurídico y político de un sistema en crisis* (October 8, 2025). https://www.hayderecho.com/2025/10/08/pulseras-antimaltrato-analisis-juridico-politico/ ↩

[6] Constitutional Court — *STC 114/1984, of November 29* (recording one's own conversation does not violate Art. 18.3 of the Constitution; admissible as documentary evidence; the intimate-sphere limit; the case's labor origin). https://hj.tribunalconstitucional.es/es/Resolucion/Show/367 · https://www.boe.es/buscar/doc.php?id=BOE-T-1984-27955 ↩

[7] Legal Today — *La grabación de las conversaciones entre particulares y su manifestación jurisprudencial* (the Supreme Court's settled doctrine on participants' recordings); Iberley — *La validez de las grabaciones privadas aportadas como prueba en el proceso penal.* https://www.legaltoday.com/practica-juridica/derecho-penal/penal/la-grabacion-de-las-conversaciones-entre-particulares-y-su-manifestacion-jurisprudencial-2014-08-14/ · https://www.iberley.es/revista/la-validez-las-grabaciones-privadas-aportadas-como-prueba-proceso-penal-1221 ↩

[8] Gerson Vidal (lawyers) — *¿Es delito grabar una conversación sin consentimiento?* (Art. 197.1: 1 to 4 years' prison and a 12-to-24-month fine; dissemination: 2 to 5 years); Noticias Jurídicas — *Voces ocultas: la delgada línea entre grabar legalmente y cometer un delito.* https://www.gersonvidal.com/blog/delito-grabar-conversacion/ · https://noticias.juridicas.com/conocimiento/tribunas/18542-voces-ocultas:-la-delgada-linea-entre-grabar-legalmente-y-cometer-un-delito-/ ↩

[9] Digital Perito — *Clonación de voz: deepfakes de audio y detección forense* (convincing clones with 3-5 minutes of audio; some systems with 30 seconds; spectrogram and prosody analysis). https://digitalperito.es/glosario/clonacion-voz/ ↩

[10] Economist & Jurist — *La prueba sintética en el proceso judicial: retos probatorios de los deepfakes* (challenging shifts the burden; the Supreme Court demands a technical, not generic, challenge; reinforced admission criteria); Rubén Blázquez (Instituto Vasco de Derecho Procesal) — *Deepfakes en el procedimiento probatorio*. https://www.economistjurist.es/noticias-juridicas/la-prueba-sintetica-en-el-proceso-judicial-retos-probatorios-de-los-deepfakes/ · https://www.institutovascodederechoprocesal.com/sites/default/files/files/Ruben%20Blazquez2.pdf ↩

[11] ICPM — *La Fiscalía señala que sólo el 0,01 % de las denuncias por violencia machista son falsas* (79 final convictions, 0.0075%). https://www.icpm.es/actualidad_noticias.asp?id=2382 ↩

[12] Efeminista — *Casi 3 millones de denuncias por violencia de género en España desde 2004 y solo un 0,0082 % son falsas* (2009-2024 series); ANAVID — *¿De dónde sale la cifra del 0,01 %?* (methodological nuance: only falsehood attested with intent). https://efeminista.com/denuncias-falsas-violencia-de-genero-espana/ · https://anavid.es/de-donde-sale-la-cifra-del-001-de-denuncias-falsas-de-violencia-de-genero-por-parte-de-la-fiscalia-del-estado-y-el-cgpj/ ↩

[13] Motorpasión — *¿Puedes usar una dashcam en España?* (AEPD criteria: recording limited to the incident, ~20 seconds before and after, preferably without sound; valid as evidence via legitimate interest and judicial protection); Adasfera — *Legalidad dashcam España: guía DGT y AEPD (2026)*. https://www.motorpasion.com/seguridad/pregunta-puedes-usar-dashcam-espana-respuesta-puedes-meterte-buen-lio-no-sabes-como · https://adasfera.com/2026/01/legalidad-dashcam-espana-dgt-multasson-legales-las-dashcams-en-espana-en-2026-guia-definitiva-dgt-y-aepd/ ↩

[14] TodoDashcam — *Legislación sobre dashcams en España y Europa* (Portugal: Law 58/2019, banned even to carry one; Germany: admission limited to the incident; France: CNIL, driver responsible for the processing). https://www.tododashcam.es/pages/blog-legislacion-dashcams-espana ↩

[15] AlertCops (State Secretariat for Security, Ministry of the Interior) — *AlertCops incorpora un botón SOS* (5 presses in under 6 seconds; location + 10 seconds of audio to the security forces; AVDG code for victims of gender violence). https://alertcops.ses.mir.es/publico/alertcops/actualidad/detalleActualidad.html?entrada=AlertCops-incorpora-un-bot-n-SOS ↩

[16] Rastreator — *Guardia Civil y UNESPA, contra el fraude* (≈400,000 fraud attempts a year; ≈€2 billion; 6 in 100 injuries; staged crashes: 3.2% of those declared, 18% of the investigated cost; GC-UNESPA Operating Procedure). https://www.rastreator.com/seguros/noticias/guardia-civil-y-unespa-luchan-contra-el-fraude-en-el-seguro.aspx ↩

[17] Infobae — *Detenidas 32 personas que fingían accidentes de tráfico para estafar al seguro* (July 2, 2025; seven staged crashes in eight months). https://www.infobae.com/espana/2025/07/02/detenidas-32-personas-que-fingian-accidentes-de-trafico-para-estafar-al-seguro-una-agencia-de-detectives-presento-la-denuncia/ ↩

[18] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; the content never leaves the device; voluntary disclosure). See chapters 13 and 15. ↩

Genomic Data: One Person's Consent Exposes Many

"...one of the most powerful and dangerous ideas in the history of science: the 'gene,' the fundamental unit of heredity."

— Siddhartha Mukherjee, The Gene: An Intimate History (2016)

On April 24, 2018, the Sacramento police arrested Joseph James DeAngelo, a seventy-two-year-old former police officer, as the "Golden State Killer," responsible for more than a dozen murders and fifty rapes in the seventies and eighties. The lead did not come from any police database. Investigators uploaded a DNA profile from the crime scene to a public genealogy database, GEDmatch, and the forensic genealogist Barbara Rae-Venter traced it through **third- and fourth-degree cousins** who had voluntarily uploaded their DNA [7]. The killer was not betrayed by his own DNA, but by **his family's genealogical tree**.

That case, at once a victory for justice and a privacy nightmare, contains this chapter's central paradox, distinct from chapter 47's — "who are you?" — and from chapter 49's — the vulnerable refugee's biometrics. Here the subject is the **consumer** who voluntarily pays for a DNA test, and the datum has a property no other has: **it is yours and your family's at once, it is irrevocable, and it betrays people who never consented to anything**. One person's consent exposes many.

The datum you can neither change nor take back

The fragility of this datum was laid bare in October 2023, when 23andMe suffered a breach. The attackers directly entered only some fourteen thousand accounts — 0.1% — but from there, through the feature that connects relatives, they exposed the data of **6.9 million users** [1]. And they did something chilling: they published lists segmented by ethnic origin — a million profiles labeled as Ashkenazi Jews, a hundred thousand as Chinese — for sale for a few dollars. Genetic data is, by construction, data of **belonging to a group**. Worst of all, it cannot be undone: unlike a password, you cannot "change your DNA" after a leak.

And in 2025 came the second blow. 23andMe, which held the data of close to **fifteen million people**, went bankrupt; within twenty-four hours traffic to its website rose 526% and close to **1.9 million people** asked to delete their DNA, on discovering that the genetics they had entrusted to a startup was now a **bankruptcy asset** to be auctioned [3]. A coalition of **28 attorneys general**, led by New York's, Letitia James, sued the company with a blunt argument: "23andMe cannot auction the genetic information of millions of people without their consent" [4]. The data ended up in the hands of a foundation of the cofounder for $305 million, and the British regulator fined the company £2.31 million — punishing, as it stressed, the **custody**, not the content [6]. And the U.S. genetic nondiscrimination law, let us remember, **does not cover life insurance** [9].

The family that never consented

Here lies the most uncomfortable truth, and it is mathematical. The researchers **Yaniv Erlich and his colleagues** showed, as we have seen, that it is enough for **2% of a population** to be in a genetic database to identify almost anyone through a distant relative; already today, close to 60% of Americans of European descent are identifiable even if **they have never taken a test** [2]. Genetic privacy, then, is not an individual right, but a **collective and family** one. As the jurist **Taner Kuru** argues, uploading your DNA does not make your cousin's DNA "public": no signature of yours can legitimize a datum that betrays children, cousins and descendants not yet born. And the same goes for consumer biometrics — the face, the iris — which cannot be changed either: under Illinois law, Facebook paid $650 million for tagging faces without consent.

The proof of existence: sealing the consent, not the genome

This is where **t2p** must be especially humble, because what it can do is valuable but limited. It does not anchor **the DNA sequence** — putting it on the chain would be catastrophic — but **the act and the terms of the consent**. The customer signs a form that fixes exactly what use they authorize — the test, the relative-matching, the research, police access; the **SHA-512 fingerprint** is computed on the device,

and only this is anchored to the **public blockchain** with an **eIDAS timestamp** and **forensic GPS**. The result is unalterable proof of *what this person consented to, when and where*. The **proof of life** binds the consent to a living, present person, and the **versioning** of each privacy policy — like GEDmatch's change in 2019 or the transfer to the foundation in 2025 — leaves a verifiable record of which consent basis actually covered each transfer. The **.asice** container is handed to the consumer, who thus keeps their **own evidentiary copy** of what they authorized, without depending on the server of a company that can go bankrupt.

The limit must be said in all its rawness, because here it is the hardest in the whole book. "Garbage in, garbage on-chain," and worse: **genetic data does not belong to whoever seals it**. First, the proof of existence **gives not privacy but traceability**: sealing the consent proves it existed, but does not prevent the data from leaking, being auctioned in a bankruptcy or being used by a life insurer; integrity is not confidentiality. Second, and gravest, the datum is **irrevocable** and the consent is structurally impossible to give in full, because **your consent seals your family's DNA**; the proof can show that *you* consented, but it does not create your third cousin's consent. Third, it does not resolve the asymmetry of the fine print: a perfectly sealed opt-in is still exactly what the user accepted without reading — the chain attests to the click, not to the click being informed. And fourth, deleting the fingerprint **does not delete the physical sample**, or the backups, or the copies already shared with researchers or attackers. **t2p** can provide **evidentiary sovereignty over the consent**, but it cannot give privacy back to a datum that, by nature, already belongs to others too and cannot be undone.

The promise and the peril

The promise, properly bounded, is real. Giving every person a sealed proof, in their own hands, of exactly what they consented to and when would make it much harder for a bankrupt company to auction data under a consent no one gave, or to rewrite the terms after the fact. In a world where the genetics you entrust to a startup can end up at auction, holding your own copy of what you authorized is a real shield against corporate repudiation.

The peril, here, is promising what cannot be delivered. Selling the proof of existence as a guarantee of genetic privacy would be a dangerous lie: the seal gives traceability, not confidentiality, and it cannot legitimize a datum that compromises the family. Truly informed consent, the ban on auctioning genomes, closing the life-insurance loophole and data minimization are problems of law and governance, not of cryptography. The proof of existence is a tool for sovereignty over consent, not a substitute for a law protecting a datum that cannot be changed.

It's time to act

For consumers: think twice before handing over a datum that also belongs to your children and cousins, and keep proof of what you authorize. For genomics and biometrics companies: seal the consent and its versioning, and require a real opt-in before any sale or transfer. For regulators: punish bad custody, as the British regulator did, and close the life-insurance loophole. And for legislators: recognize that genetic privacy is familial and collective, and that no individual signature can sell the DNA of an entire lineage.

A killer fell because a relative had uploaded their DNA; millions of people saw their genome up for sale after a company went bankrupt. Giving every person proof of what they have consented to — never forgetting that the seal does not restore privacy to an irrevocable, shared datum — is, on this front that touches what is most intimate in all of us, time2act.

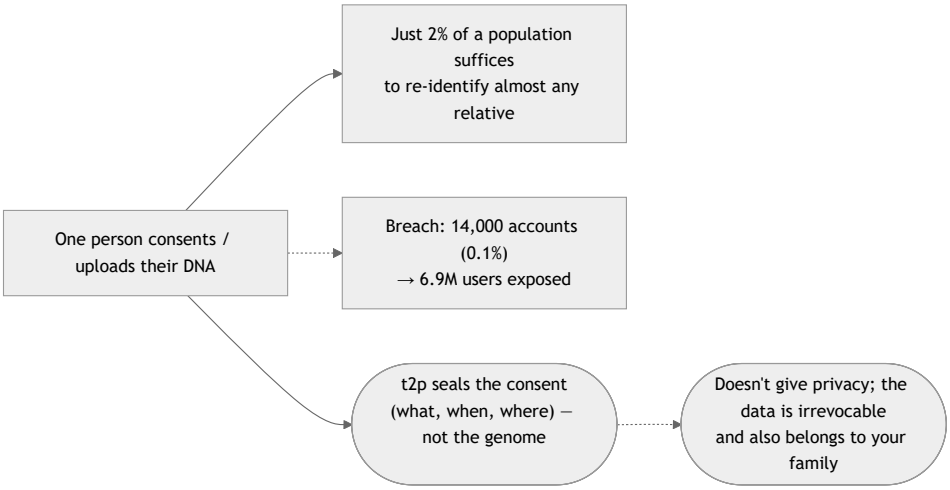


Figure 76.1 — Two truths of genetic data. On the left, the family effect: it is enough for 2% of a population to be in a database to reidentify almost anyone through a distant relative, and close to 60% of those of European descent are already identifiable even if they never took a test (Erlich et al., 2018). On the right, the breach's multiplier effect: the attack on 14,000 accounts (0.1%) exposed 6.9 million users through the relative-matching feature. t2p can seal the act of consent — what, when and where — but it cannot restore privacy to an irrevocable datum that also belongs to your family. Sources: Science (Erlich et al.); TechCrunch; New York attorney general.

---

[1] TechCrunch — *23andMe confirms hackers stole ancestry data on 6.9 million users* (2023; the lists segmented by ethnic origin). https://techcrunch.com/2023/12/04/23andme-confirms-hackers-stole-ancestry-data-on-6-9-million-users/ ↩

[2] Yaniv Erlich, Tal Shor, Itsik Pe'er and Shai Carmi (Science) — *Identity inference of genomic data using long-range familial searches* (2018; 2% of the population is enough). https://www.science.org/doi/10.1126/science.aau4832 ↩

[3] CNBC — *23andMe bankruptcy: selling and deleting DNA* (2025; +526% traffic; ~1.9 million deletion requests). https://www.cnbc.com/2025/03/30/23andme-bankruptcy-selling-deleting-dna-genetic-testing.html ↩

[4] New York Attorney General (Letitia James) — *AG James sues 23andMe to protect genetic data* (a coalition of 28 attorneys general). https://ag.ny.gov/press-release/2025/attorney-general-james-sues-23andme-protect-new-yorkers-genetic-data ↩

[5] Bio-IT World — *Anne Wojcicki buys back 23andMe for $305M* (Regeneron, $256 million → TTAM, $305 million; closing on 14-Jul-2025). https://www.bio-itworld.com/news/2025/07/02/anne-wojcicki-buys-back-23andme-for-305m-promises-data-security

[6] Information Commissioner's Office (United Kingdom) — *ICO fines 23andMe £2.31 million* (155,592 UK residents; the penalty punishes the custody). https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2025/06/23andme-fined-for-failing-to-protect-uk-users-genetic-data/ ↩

[7] NPR — *In hunt for the Golden State Killer, investigators uploaded his DNA to a genealogy site* (the DNA was uploaded by a relative). https://www.npr.org/sections/thetwo-way/2018/04/27/606624218/in-hunt-for-golden-state-killer-investigators-uploaded-his-dna-to-genealogy-site ↩

[8] The Legal Genealogist — *GEDmatch reverses course* (opt-out by default; only ~14% opted in). https://www.legalgenealogist.com/2019/05/19/gedmatch-reverses-course/

[9] National Human Genome Research Institute — *GINA: Genetic Information Nondiscrimination Act (2008)* (does not cover life, disability or long-term care insurance). https://www.genome.gov/27026050/president-bush-signs-the-genetic-information-nondiscrimination-act-of-2008 ↩

[10] Taner Kuru (Computer Law & Security Review) — *Investigative genetic genealogy in Europe* (2025; uploading your DNA does not make your relatives' "public"). https://www.sciencedirect.com/science/article/pii/S0267364925000019

Criminal Justice: The Evidence a Court Must Be Able to Believe

"It is better that ten guilty persons escape than that one innocent suffer."

— William Blackstone, Commentaries on the Laws of England, IV, 27 (1769)

On January 24, 2017, Baltimore police officer Richard Pinheiro walked into a litter-strewn yard, placed a bag of capsules inside a soup can, stepped out into the alley and *then* switched on his body camera to record the "discovery" of the drugs. What he did not know is that the device automatically preserves the **thirty seconds before activation, without sound**. Those thirty silent seconds gave him away: he was convicted of fabricating evidence, and the drug case collapsed [4]. The very technology designed to watch the citizen had recorded the watchman.

That case contains this chapter's two faces, distinct from chapter 20 — which dealt with evidence of atrocities before an international tribunal. Here the terrain is **ordinary criminal justice**: an officer's body camera, the audio in a divorce, the video in a national trial. And it poses the question that today is cracking courts all over the world: in an era when any sound or image can be convincingly faked, **how can a court know which evidence to believe?**

The liar's dividend inside the courtroom

The threat runs in two directions. The first is the fake that deceives: an AI-fabricated audio or video presented as authentic. In a British family case, a mother submitted an audio in which the father supposedly made violent threats; the experts showed, analyzing the metadata, that it had been edited after the call to insert words the father never said [9].

The second direction is more insidious and has a name of its own. The jurists Robert Chesney and Danielle Citron, in their reference article *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security* (2019), christened it the "liar's dividend": as everyone comes to assume that everything can be faked, some will dodge accountability by denouncing an authentic recording as a deepfake [6].

And it has already reached the courts. In the litigation over the death of an engineer at the wheel of a Tesla, the company's lawyers went so far as to argue that the videos in which Elon Musk touts the safety of his system "could be *deepfakes*." Judge **Evette Pennypacker** called it "deeply troubling": by that logic, any public figure could say whatever they pleased and then deny it [1]. And in the Kyle Rittenhouse trial, the defense got the judge to doubt whether zooming in on a video on a tablet — a simple pixel interpolation — "altered it with artificial intelligence," and gave the prosecution only twenty minutes to produce an Apple engineer; the enlarged image was never shown to the jury [2]. When cloning a convincing voice takes only a few seconds of audio and, according to a 2025 study, only 0.1% of people reliably tell real from fake, doubt becomes a weapon.

Who controls the archive

There is a third problem, more prosaic but just as corrosive: **who holds the evidence**. The body camera was sold as a transparency tool, but it has often become a tool controlled by the police themselves. A **ProPublica** investigation found that, of seventy-nine police killings on camera in the United States, only 42% had the footage published a year and a half later, and that in roughly one in five cases the department falsely told the oversight body that the footage **did not exist** [5]. In England and Wales, a BBC investigation documented more than a hundred and fifty reports of camera misuse — officers switching them off, deleting or sharing footage — facts which, according to one senior officer, "go to the heart of what undermines trust in the police" [10]. As The Sedona Conference, the evidence-law think tank, sums up: in the digital era evidence is not "stored," it is **held in custody with traceability** — without a provable chain, the most valuable datum is impeachable.

The proof of existence: sealing in the second of recording

This is where **t2p** attacks the problem at its exact origin. At the instant of capture, the body camera or the witness's phone computes **on the device** the SHA-512 fingerprint of the file — which never leaves; only

the fingerprint travels — anchors it to the **public blockchain** and attaches an **RFC 3161 timestamp** conforming to eIDAS, a **forensic GPS** fix and an **.asice** container. The seal is not applied *afterward*, in the evidence room: it is applied **in the second of recording**, closing the manipulation window.

The effect on the two threats is direct. Against later editing or deletion — the Baltimore case or the BBC's hundred and fifty reports — any cut or "loss" of a sealed clip would fail verification. And against the liar's dividend, faced with "this video is a *deepfake*," the independent anchor proves that **that exact file existed intact before such an instant**, shifting the burden of doubt onto whoever challenges it. All of this fits existing law: U.S. federal rule 902(14) already admits the **hash** as a route to self-authentication of digital copies, and the European eIDAS regulation gives the qualified timestamp a **legal presumption** of integrity and date. And the .asice container lets both the defense and the prosecution verify integrity **without the app and without trusting t2p**, using open tools: a guarantee of equality of arms, because the defense can independently audit the prosecution's evidence.

The limit must be said in all its rawness, because the Baltimore case is the devastating example. A body camera sealed **at the moment** of recording Pinheiro planting the drugs would have certified, with an unalterable seal, a fraudulent scene. "Garbage in, garbage on-chain": if the officer acts in bad faith in front of the camera, the chain certifies the lie with a certain date. The seal proves integrity, instant, place and person, **not the truth of the scene**. And there are limits specific to this front. First, the proof only works if the seal is applied **at origin**: it does not rescue the historical unsealed recordings, which today are 99.9% of the real evidence in the courts. Second, **it does not force the police to switch the camera on**: ProPublica's problem is not the clip's integrity but its *nonexistence*, and no seal creates recordings no one made. Third, against the liar's dividend, the anchor only helps if the evidence **was** sealed; for the ocean of unsealed material, the doubt remains intact, and a system that vouches only for sealed material could, perversely, **increase** suspicion of everything else. And fourth, final admissibility is decided by the judge under national rules, not by the blockchain: the hash self-authenticates the *file*, not its relevance or the lawfulness of how it was obtained.

The promise and the peril

The promise is to restore the base of facts without which no trial is possible. Evidence sealed at origin would give the court a way to tell the intact file from the fabricated one, would armor the body camera's chain of custody and would break the liar's dividend for protected material. The very body that writes the evidence rules in the United States, the Advisory Committee chaired by Professor Daniel Capra, is already debating new rules for machine-generated evidence, acknowledging that neither the tools nor the judges yet know how to spot deepfakes reliably. The seal of origin is exactly the authentication reinforcement this debate is looking for.

The peril calls for the nuances we have stated, and here they are crucial. The first, do not confuse integrity with truth: a sealed camera that records a setup certifies the setup. The second, the seal neither forces the camera on nor rescues the old material. And the third, the subtlest: if only what is sealed is trusted, everything else may end up unjustly distrusted. The proof of existence is a powerful reinforcement of evidentiary authentication, not a lie detector or a substitute for human judgment.

It's time to act

For **police forces**: seal every recording in the second of capture and make camera activation mandatory, so that the clip's integrity and existence do not rest in the hands of the party under investigation. For **courts and rule-writers** (the Advisory Committee, legislators): recognize the seal of origin as a route to self-authentication and adapt the rules to machine-generated evidence. For the **defense and the prosecution**: demand independently verifiable containers, to guarantee equality of arms. And for **citizens**: a way to prove that what they recorded is exactly what happened.

An officer was betrayed by his own camera, but a sealed camera could just as well have certified his lie had he recorded it with a poker face. Giving every piece of evidence a fingerprint sealed at the moment of birth — never forgetting that the seal does not make the scene true — is, on this front on which justice depends, time2act.

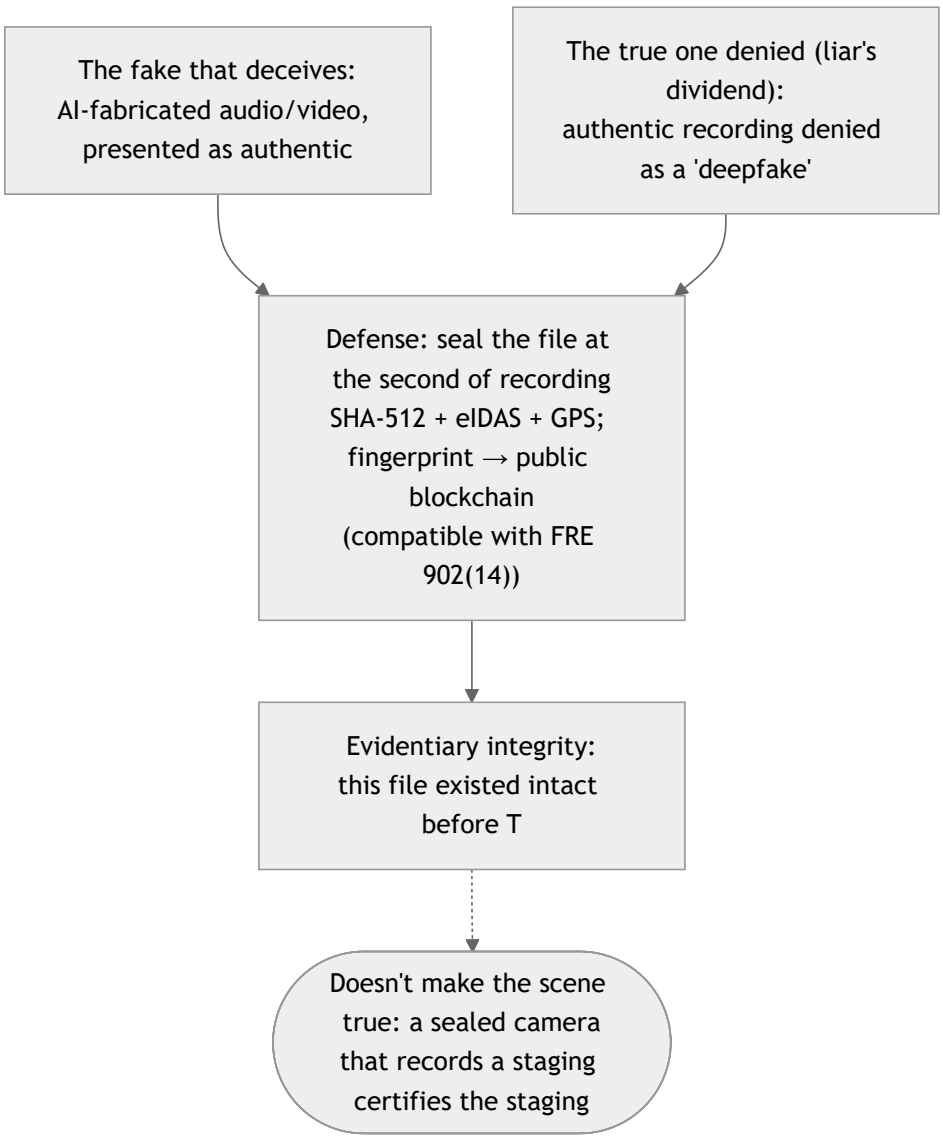


Figure 77.1 — Digital evidence suffers two opposite attacks. On the left, the fake that deceives: an AI-fabricated audio or video presented as authentic (like the edited audio in a British family case). On the right, the liar's dividend: an authentic recording denied as a "deepfake" to dodge accountability (as in Huang v. Tesla, where it was argued that the Musk videos "could be deepfakes"). The same defense works for both: sealing the file in the second of recording with a SHA-512 fingerprint, an eIDAS seal and GPS (compatible with rule 902(14): the hash self-authenticates copies). Limit: the seal proves integrity and moment, not the truth of the scene — a sealed camera that records a setup certifies the setup. Sources: Chesney and Citron (2019); ProPublica; the Tesla and Baltimore cases.

---

[1] Fortune — *Tesla lawyers argue recordings of Musk touting Autopilot could be deepfakes* (Huang v. Tesla; Judge Evette Pennypacker). https://fortune.com/2023/04/27/elon-musk-lawyers-argue-recordings-of-him-touting-tesla-autopilot-safety-could-be-deepfakes/ ↩

[2] 9to5Mac — *Rittenhouse lawyer's pinch-to-zoom argument* (the burden shifted onto the prosecution; 20 minutes for an Apple expert). https://9to5mac.com/2021/11/11/rittenhouse-lawyer-pinch-to-zoom/ ↩

[3] NBC News — *Washington state judge blocks use of AI-enhanced video evidence* (State v. Puloka; Judge Leroy McCullough; 29-Mar-2024). https://www.nbcnews.com/news/us-news/washington-state-judge-blocks-use-ai-enhanced-video-evidence-rcna141932

[4] NPR — *Baltimore police caught planting drugs in body-cam footage* (Officer Pinheiro; the 30-second soundless buffer). https://www.npr.org/sections/thetwo-way/2017/07/20/538279258/baltimore-police-caught-planting-drugs-in-body-cam-footage-public-defender-says ↩

[5] ProPublica — *How police have undermined the promise of body cameras* (only 42% of 79 killings; a false "does not exist" in 1 in 5). https://www.propublica.org/article/how-police-undermined-promise-body-cameras ↩

[6] Robert Chesney and Danielle Keats Citron — *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security* (107 California Law Review 1753, 2019; the "liar's dividend"). https://www.californialawreview.org/print/deep-fakes-a-looming-challenge-for-privacy-democracy-and-national-security ↩

[7] U.S. Courts — *Advisory Committee on Evidence Rules: proposed Rule 707 (machine-generated evidence) and 901(c)* (comment period until 16-Feb-2026). https://www.uscourts.gov/sites/default/files/document/2025-11_evidence_rules_commitee_agenda_book_final.pdf

[8] American Bar Association — *New rules for self-authenticating electronic evidence (FRE 902(13)/(14))* (the hash value self-authenticates digital copies). https://www.americanbar.org/groups/litigation/resources/newsletters/trial-evidence/new-rules-self-authenticating-electronic-evidence/

[9] CYFOR / Expatriate Law — *Deepfake audio evidence used in a UK court* (audio edited after the call to insert threats). https://cyfor.co.uk/deepfake-audio-evidence-used-in-uk-court-to-discredit-father/ ↩

[10] BBC (via LBC) — *150 reports of police officers misusing body cameras* (England and Wales; Jim Colwell, NPCC). https://www.lbc.co.uk/article/150-reports-police-officers-misusing-body-cameras-investigation-england-wales-DWzgZZ_2/ ↩

Medical Records: Who Touched the Record, and When

"It may seem a strange principle to enunciate as the very first requirement in a hospital that it should do the sick no harm."

— Florence Nightingale, Notes on Hospitals (1863)

In Chicago, a mother sued Rush University Medical Center on behalf of her son, born with a brain injury from lack of oxygen during delivery. To prove what had happened minute by minute in the room, she needed one specific thing: the **audit trail** of the electronic record, the trace showing who wrote, modified or deleted each entry and when. For more than three years, the hospital handed over a simple access log instead of the complete trail and made claims about its ability to produce it that the judge deemed **false**. In 2022, Judge James O'Hara imposed a sanction equivalent to a **default judgment**: liability was fixed; all that remained was to put a figure on it [1]. The irony is cruel: the document being litigated was, precisely, the one that was supposed to tell the truth of what happened.

That withheld trail contains this chapter's question, distinct from chapter 18's — which deals with medical imaging and prescriptions. Here we are not speaking of the clinical *content*, but of the **governance of the record**: **who accessed a chart, who altered it, when, and did the patient consent?** Because an electronic record, however sophisticated it looks, is a record that can be rewritten, backdated and hidden.

The record that can be rewritten

Chart manipulation is a documented and technically simple fraud. Forensic examiners have found cases where a shift's notes were written **after the patient's death and backdated**, made from a non-clinical workstation hours after the cardiac arrest: what on paper looked contemporaneous was, in the metadata, later. That is why the American drug agency's rule is categorical: a change in an electronic record **cannot obscure the previous information** [3]. At a larger scale, the fraud takes the form of **copy-paste**: the American health system's inspection office found that only a quarter of hospitals had policies on this function, and that stays billed at the highest severity level grew almost 20% in five years, coming to represent **almost half of all Medicare spending** on hospitalizations [4][5]. And there is the snooping: prying into celebrities' charts cost the UCLA hospital an **$865,500** fine in 2011, and another former employee of the same center, **Huping Zhou**, who had accessed records **323 times**, became **the first person jailed in the United States for a violation of the health privacy law** [6][7]. The audit trail reveals all these frauds — but, as the Prieto case showed, it can also be withheld.

Consent as a receipt, not a checkbox

There is a second dimension: **consent**. Today the patient's consent is usually a checkbox lost in a form, impossible to audit and easy to fabricate after the fact. The **Kantara** initiative, in its *consent receipt* specification — the basis of the international standard ISO 27560 — proposes a simple, powerful idea: that consent materialize in a **receipt readable by humans and processable by machines**, like a purchase receipt, recording who consented, for what purpose and under which policy. Consent thus becomes **auditable and revocable**, not a forgotten checkbox. And U.S. law already pushes in this direction: the Cures Act obliges providers to give patients access to their health information and prohibits "information blocking" — exactly what the hospital did in the Prieto case.

The proof of existence: sealing the event, not the image

This is where **t2p** contributes a layer of integrity to the record and the consent. What is sealed **is not the image or the clinical datum** — which stays on the hospital's server — but the **event**. Every time an entry is created, modified or deleted, the device computes the **SHA-512 fingerprint** of the pair formed by the entry's content and its audit metadata — user identifier, role, time, action type — and this fingerprint is anchored to the **public blockchain** with an **eIDAS timestamp**. A public, unalterable proof is generated that *that* entry existed with *that* content at *this* instant. The same applies to the **consent receipt** — proving when it was granted and, above all, when it was **revoked**, preventing

fabricated retroactive consent — and to the **access receipt**, which creates a record of who has consulted the chart that can no longer vanish, as in the UCLA case.

The effect is direct. The **proof of life** binds the signing of a note to a real, present clinician, not to a stolen session; the **forensic GPS** corroborates that the entry was made from a plausible clinical location, not a remote workstation; the **timestamp** fixes the instant independently of the record system's manipulable clock; and the **.asice** container packages the entry, its consent receipt and the anchoring proof so that an expert can verify them **without access to the hospital's system**. Backdating, overwriting and the selective filtering of the trail in litigation become detectable.

The limit must be said in all its rawness, because here it is fundamental. "Garbage in, garbage on-chain": the sealing proves the record's **temporal and authorship integrity**, not its **clinical truth**. If a doctor writes in real time a false observation — "patient stable" when they were not — and seals it instantly, t2p will certify forever that that lie existed at that moment and by that hand, and will make it **harder** to retract, not easier. The proof detects neither negligence nor clinical error; it only freezes what was written. Moreover, it does not replace the record system's native audit trail: if the hospital does not generate or seal an event, there is nothing to anchor, and an access outside the sealed flow is invisible to the chain. The consent receipt proves the patient signed, not that they understood or that the consent was free. And the record of who looks can reveal sensitive **patterns**, something that demands governance. The proof of existence gives the record integrity, not veracity to its content.

The promise and the peril

The promise is to restore trust in the document on which life and justice depend. A chart with every event sealed at origin would make it impossible to backdate a note after a bad outcome, to hide the trail in litigation as in the Prieto case, or to erase a snooper's tracks. And it would give the patient real control over their consent, with a receipt proving what they authorized and when they revoked it. In a system where the record can be rewritten, making the event unalterable protects both the patient and the honest clinician, who could prove their note is exactly the one they wrote.

The peril calls for the nuances we have stated. The first, do not confuse integrity with truth: a false note, sealed, is still false and now, moreover, immutable. The second, the seal neither replaces the native log nor rescues what was not sealed. And the third, the privacy of the access trail itself, which must be protected so it does not become a tool for surveilling the clinician. The proof of existence is a powerful layer of record governance, not a certificate of good practice or a substitute for the expert judgment that knows how to read an audit trail.

It's time to act

For **hospitals and record-system makers** (Epic, Cerner): seal every write, access and consent event at origin, so the trail can be neither backdated nor hidden. For **regulators** (the health, drug and interoperability agencies): recognize the seal of origin as compliance with the audit and integrity controls the law already requires. For **patients**: a consent receipt in their own hands proving what they authorized and when they revoked it. And for **courts**: treat the sealed audit trail as the reliable evidence that was missing in the Prieto case.

A mother fought for three years for a record that was supposed to tell the truth of her son's birth, and the hospital hid it. Giving every event in a chart a sealed proof of who touched it and when — never forgetting that the seal does not make what is written in it true — is, on this front that touches health and justice, time2act.

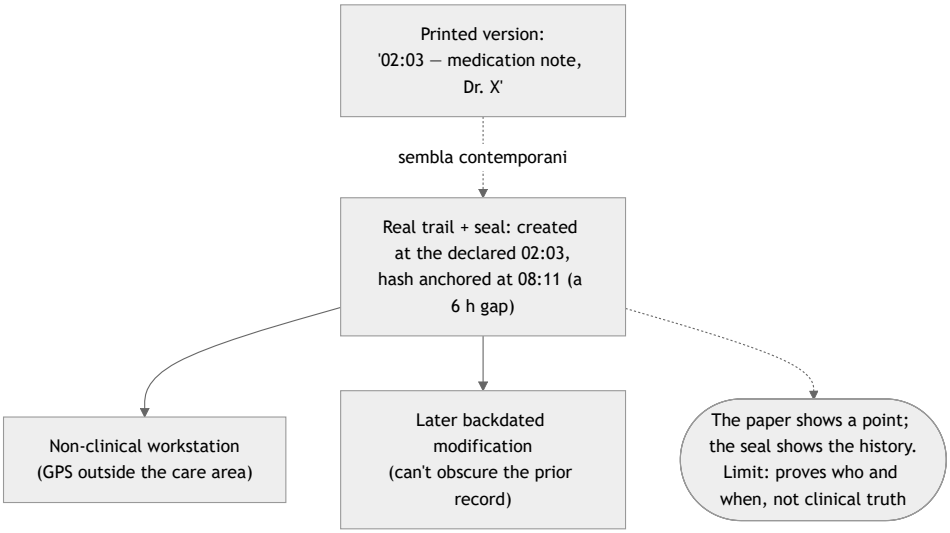


Figure 78.1 — A clinical note seen in two lanes. Above, the printed version: a single point, "02:03 — medication note, Dr. X." Below, the real audit trail with the t2p seal: the creation event declared at 02:03 but with the fingerprint anchored to the public blockchain at 08:11 — a six-hour gap — made from a non-clinical workstation (forensic GPS outside the care area), with a later backdated modification that cannot obscure the previous information (21 CFR 11.10(e)). The paper shows a point; the seal shows the story. Limit: the seal proves who and when, not the clinical truth of what was written. Sources: HIPAA and FDA (audit controls); Prieto v. Rush; forensic analysis of records.

---

[1] Enigma Forensics — *Judge sanctions hospital for not producing complete EMR audit trail* (Prieto v. Rush University Medical Center; Judge James O'Hara; "information blocking"). https://enigmaforensics.com/blog/judges-order-sanctions-hospital-not-producing-complete-emr-audit-trail/ ↩

[2] Cornell Law (Legal Information Institute) — *45 CFR §164.312: HIPAA technical safeguards* (mandatory audit controls; ePHI integrity). https://www.law.cornell.edu/cfr/text/45/164.312

[3] Cornell Law (Legal Information Institute) — *21 CFR §11.10: FDA electronic records* (timestamped audit trails; a change cannot obscure the previous record). https://www.law.cornell.edu/cfr/text/21/11.10 ↩

[4] HHS Office of Inspector General — *Not All Recommended Fraud Safeguards Have Been Implemented in Hospital EHR Technology* (OEI-01-11-00570, 2013; only ~25% with copy-paste policies). https://oig.hhs.gov/oei/reports/oei-01-11-00570.asp ↩

[5] HHS Office of Inspector General — *Trend Toward More Expensive Inpatient Hospital Stays in Medicare* (OEI-02-18-00380, 2021; almost half the spending at the highest severity level). https://oig.hhs.gov/oei/reports/OEI-02-18-00380.pdf ↩

[6] ProPublica / HHS OCR — *UCLA Health System pays $865,500 to settle celebrity-privacy allegations* (2011; the snooping in the charts). https://www.propublica.org/article/ucla-health-system-pays-865000-to-settle-celebrity-privacy-allegations ↩

[7] U.S. Department of Justice — *Former UCLA healthcare worker sentenced (Huping Zhou)* (323 accesses; the first prison sentence for a HIPAA violation). https://www.justice.gov/archive/usao/cac/Pressroom/pr2010/004.html ↩

[8] Federal Register / ONC — *21st Century Cures Act: interoperability and information blocking* (2020; patient access to their health information). https://www.federalregister.gov/documents/2020/05/01/2020-07419/21st-century-cures-act-interoperability-information-blocking-and-the-onc-health-it-certification

[9] Kantara Initiative — *Consent Receipt Specification v1.1* (basis of the ISO/IEC TS 27560:2023 standard). https://kantara.atlassian.net/wiki/spaces/archive/pages/3508790/Consent+Receipt+Specification

[10] Dean F. Sittig and Adam Wright (Applied Clinical Informatics) — *Identifying a Clinical Informatics or EHR Expert Witness* (2023; the log is evidence that must be interpreted). https://pmc.ncbi.nlm.nih.gov/articles/PMC10033222/

AI Models: What They Were Trained On, and What We Actually Run

"We had better be quite sure that the purpose put into the machine is the purpose which we really desire."

— Norbert Wiener, "Some Moral and Technical Consequences of Automation," Science (1960)

In December 2023, a team from the **Stanford Internet Observatory** looked inside something no one had truly audited for years: **LAION-5B**, the dataset of five billion image-text pairs on which Stable Diffusion and much of AI image generation had been trained. They found thousands of links to child sexual abuse material — more than **3,200 suspicious cases**, a portion confirmed by third parties — hidden among the web-scraped data [1]. LAION pulled the dataset off the network, and the report recommended **deprecating and discontinuing** the models trained on it; months later it published a cleaned version, Re-LAION-5B, from which it removed more than two thousand links [1][2]. The underlying problem was not just the content: it was that **no one — not even its creators — could say with certainty what was inside** until someone from outside audited it. A model that millions of people had already downloaded had been trained on a pile of data that turned out to be, in part, illegal, and no reliable record existed to say so.

This is this chapter's question, distinct from chapter 40's — which deals with the rights of whoever creates a work. Here we are not speaking of the content's authorship, but of the **provenance of the model itself**: what data it was trained on, who built it, when, and — no less important — whether the **weights** you downloaded and run are exactly the ones someone evaluated. Because an AI model, however impressive, is an opaque artifact: a weights file and a documentation card that can be rewritten, backdated and, as we shall see, even infected.

The model no one can audit

The opacity has three faces. The first is the **data**: large models are trained on corpora scraped from the whole internet, so enormous that their composition is known only by approximation; the LAION case showed that not even the creators know for sure what they contain. The second is the **documentation**: since 2019, the *model cards* proposed by Mitchell and colleagues have standardized how to declare a model's intended use, performance and limits [9]; but a model card is a voluntary document living in a repository the provider controls, and which it can edit — or soften — after an incident. And the third, the most unsettling, is the **integrity of the weights**. In 2024, JFrog researchers found **a hundred malicious models** posted on Hugging Face: exploiting the *pickle* serialization format — the standard for packaging models — they hid code that executed on loading the file and opened a **backdoor** on the victim's machine [3][4]. The consequence is devastating: the model you download may **not be** the one the laboratory evaluated and declared safe. Between the evaluation and your GPU there is a chain of custody that today no one seals.

The law demands paper; the anchor is missing

The law has begun to respond. The European AI Regulation obliges, since August 2, 2025, providers of general-purpose models to maintain technical documentation — architecture, number of parameters, compute resources, description of the data — and to publish a "sufficiently detailed summary" of the training content following a template from the AI Office [5][6]. Noncompliance is expensive: up to €15 million or 3% of worldwide turnover, and from August 2, 2026 the Office will be able to verify it; from the same date, Article 50 requires machine-readable marking of AI-generated content [5][6]. The industry has proposed an AI-BOM, a model's bill of materials like the one that already exists for software. The CycloneDX standard added the ML-BOM — datasets, models and configurations with their provenance — (as we have seen) and OWASP's AIBOM project is pushing its standardization [7][8]. But here we hit the same wall: a BOM, like a model card or a training summary, is only as reliable as the will of whoever signs it; a compliance PDF is just a PDF, and nothing stops it from being touched up tomorrow.

Why the manifest alone does not save you

It must be said bluntly: publishing a training summary or a *hash* of the weights in a README file proves nothing if that file can be changed without leaving a trace. A provider that, after an incident like LAION's, quietly edited its model card to remove an awkward source would leave no mark of having done so. A weights hash without an **independent certain date** does not distinguish the original from the version touched up after the fact. And an AI-BOM stored in the same repository controlled by the party who must account for it is, in a dispute, one side's paper. What is missing is not the document: it is **an unalterable anchor fixing who declared what, and when**.

The proof of existence applied

This is where **t2p** supplies the missing layer, and it does so on the digital artifact, not on a photo of the terrain. On publishing a model, the **SHA-512 fingerprint** of the weights file and of the **AI-BOM manifest** — the training data summary, the evaluation and safety report, the licenses — are computed, and both are anchored with an **RFC 3161 timestamp conforming to eIDAS** to the public blockchain, inside an **.asice** container verifiable offline. This generates public, unalterable proof that on *that* day the provider declared *that* training summary and *that* evaluation result, and that *these* weights — by their fingerprint — are the ones evaluated.

The effect is direct on the three opacities. On the **data**, the training summary gains a certain date: if tomorrow the model must be deprecated like LAION's, no one can pretend that what was declared was never declared. On the **documentation**, later silent editing becomes detectable, because the auditor can compare the current card against the version sealed on day zero. And on the **weights**, anyone deploying the model recomputes the fingerprint and compares it with the anchor: a *nullifAI*-style substitution **fails the check** before reaching production. It is, on the model, the complement of what chapters 13 and 15 propose for the output: the C2PA mark and SynthID seal **what a model produces** — Google has already marked more than a hundred billion pieces of content with SynthID [10]; t2p seals **what a model is**.

Now the limit, in all its rawness, because here it is essential. "Garbage in, garbage on-chain": the seal proves that the provider **committed** to a given training summary and that the weights matched a fingerprint at an instant; **it does not prove the summary is true** — that the data really was only what was declared — nor that the model is safe, nor does it detect a skillful poisoning of the dataset. It binds a claim to a date and an author; it does not verify its content. At Stanford, the illegal material had to be found by **humans** auditing; no seal would have done it on its own. The proof of existence does not replace dataset auditing, red-teaming or watermarking: it makes them **accountable**, because it prevents denying afterward what was declared before.

The promise and the peril

The promise is an AI ecosystem where lineage is **verifiable**: being able to prove with what data, what version and what evaluation a model was built, and that the weights running in production are the audited ones, with accountability that survives the incidents. In an industry where a weights file travels from hand to hand without any receipt, making the origin unalterable protects both the honest deployer and the public.

The peril calls for the usual nuances. First, remember that **the seal does not prove the truth of the training summary**; a false, sealed document is still false — and now, moreover, immutable — and it breeds a false reassurance ("it is anchored, therefore it is safe") that corresponds to nothing. Second, **concentration**: if only the big laboratories can seal and document, the standard may expel open models instead of strengthening them, consolidating market power that could limit diversity and innovation. Third, the tension between publishing the training summary as proof of governance and preserving legitimate trade secrets; the necessary transparency must be balanced against the protection of intellectual property and competitive advantage. The proof of existence is a layer of model governance, not a certificate of goodness or a substitute for the expert judgment that knows how to audit a corpus.

Moreover, sealing a model's provenance does not resolve the fundamental difficulty of understanding the decision criteria generated inside deep neural networks. As Cathy O'Neil points out, these systems are "opaque and invisible except to the priests of the sector: the mathematicians and the engineers" [11]. This creates a new form of concentrated power: only those with access to the internal expertise can decipher or manipulate the underlying logic, while the public and even the regulators are left with a seal of authenticity without knowing *why* the model acts as it does. In parallel, the legal framework that should

curb this asymmetry is retreating: as we have seen, the European Union has postponed the obligations for high-risk AI from 2026 to the 2027-2028 period [12], and the state of Colorado has entirely repealed its AI law, eliminating the duty not to discriminate [12]. This regulatory retreat favors an environment where oversight becomes harder and the power of those who control the models intensifies without any effective counterweight.

An unexpected reinforcement comes from public AIs themselves, since many produce intrinsically auditable records. The **Policy Synth** tool, designed by Beth Noveck, generates a detailed record — citation by citation — of all its agents' actions, and each output comes with its documentary origin [13]. This means that, beyond the blockchain seal guaranteeing immutability, a traceability layer already exists: any auditor can follow the thread of decisions back to the original source without depending on a single provider. At the municipal scale, cities like Helsinki and Amsterdam published the first algorithmic inventories (2020), exposing the training data and the human oversight mechanisms for each system [14]. Anchoring these logs and inventories with a neutral proof of existence makes them **non-repudiable** and verifiable by any actor, removing the need to "trust the administration's website" and adding a new level of decentralized trust to the fabric of algorithmic governance.

It's time to act

For **model providers** (OpenAI, Google, Meta, Mistral, the open-weights laboratories): seal the AI-BOM on publication — fingerprint of the weights, training summary and evaluation report — so the lineage has a certain date. For **deployers and companies** that integrate models: verify that the weights you run match the sealed fingerprint before deploying them, and demand an anchored AI-BOM from the provider as a purchasing requirement. For **model repositories** (Hugging Face): anchor each model's fingerprint on upload, so a malicious substitution is detectable. And for **auditors and regulators** (the AI Office): recognize the sealed AI-BOM as evidence of the documentation Article 53 already requires, and treat the certain date as the difference between a declaration and an alibi fabricated afterward.

At Stanford, a group of researchers had to audit from outside a model already running all over the world to discover what it contained. Giving every model a sealed proof of what it was trained on, who made it and which weights are the right ones — never forgetting that the seal does not make what is declared in it true — is, on this front where the machine learns from us, **time2act**.

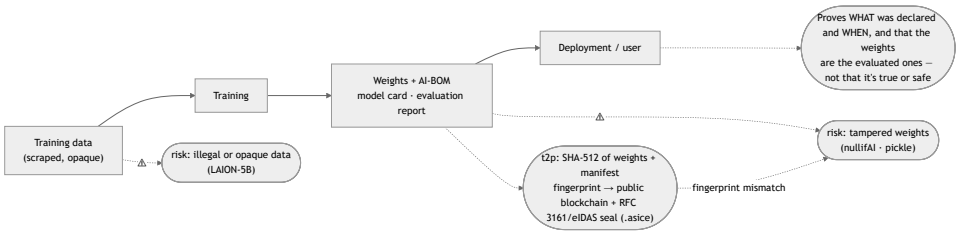


*Figure 79.1 — A model's life cycle in four links: training data → training → weights + AI-BOM (model card, evaluation report) → deployment. t2p anchors the SHA-512 fingerprint of the weights and the manifest to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). Two risk points: illegal or opaque data at the origin (LAION-5B) and weights manipulated between the laboratory and the user (*nullifAI *backdoor via* pickle*), which the fingerprint check detects. Limit: the seal proves what was declared and when, and that the weights are the evaluated ones; it does not prove the declaration is true or the model safe. Sources: Stanford Internet Observatory (LAION-5B); JFrog (malicious models); European AI Regulation, Arts. 53 and 50; CycloneDX ML-BOM.*

---

[1] Stanford Internet Observatory / 404 Media — *Largest Dataset Powering AI Images Removed After Discovery of Child Sexual Abuse Material* (December 2023; 3,226 suspicious cases in LAION-5B; recommendation to deprecate the models). https://www.404media.co/laion-datasets-removed-stanford-csam-child-abuse/ ↩

[2] LAION — *Releasing Re-LAION-5B: transparent iteration on LAION-5B with additional safety fixes* (August 2024; 2,236 links removed). https://laion.ai/blog/relaion-5b/ ↩

[3] JFrog Security Research — *Data Scientists Targeted by Malicious Hugging Face ML Models with Silent Backdoor* (2024; the *nullifAI* technique, *pickle* abuse, backdoor). https://jfrog.com/blog/data-scientists-targeted-by-malicious-hugging-face-ml-models-with-silent-backdoor/ ↩

[4] The Hacker News — *Over 100 Malicious AI/ML Models Found on Hugging Face Platform* (March 2024). https://thehackernews.com/2024/03/over-100-malicious-aiml-models-found-on.html ↩

[5] EU Artificial Intelligence Act — *Article 53: Obligations for providers of general-purpose AI models* (Annex XI technical documentation; training content summary; penalties of up to €15 million or 3%). https://artificialintelligenceact.eu/article/53/ ↩

[6] European Commission — *General-Purpose AI Models in the AI Act: Questions & Answers* (the training summary template; application from August 2, 2025; Article 50 marking from August 2, 2026). https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers ↩

[7] CycloneDX (OWASP) — *Machine Learning Bill of Materials (AI/ML-BOM)* (datasets, models and configurations with provenance; since v1.5, 2023). https://cyclonedx.org/capabilities/mlbom/ ↩

[8] *Operationalising AI bills of materials for verifiable AI provenance and lifecycle assurance* (Frontiers in Computer Science, 2026). https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2026.1735919/full ↩

[9] Margaret Mitchell et al. — *Model Cards for Model Reporting* (2019). https://arxiv.org/abs/1810.03993 ↩

[10] Google / C2PA — *How Google and the C2PA are increasing transparency for gen AI content* (SynthID: more than 100 billion pieces of content marked according to Google I/O 2026 (May 2026); content credentials). https://blog.google/innovation-and-ai/products/google-gen-ai-content-transparency-c2pa/ ↩

[11] Cathy O'Neil — *Armas de destrucción matemática* (Capitán Swing, 2018). ↩

[12] Council of the EU / Pinsent Masons (2026). https://www.pinsentmasons.com/out-law/news/rules-high-risk-ai-delayed-under-eu-omnibus-deal ↩

[13] Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 6. ↩

[14] VentureBeat (2020). https://venturebeat.com/ai/amsterdam-and-helsinki-launch-algorithm-registries-to-bring-transparency-to-public-deployments-of-ai ↩

AI Agents: The Verifiable Mandate

"Since we do not now have any ways of making computers wise, we ought not now to give computers tasks that demand wisdom."

— Joseph Weizenbaum, Computer Power and Human Reason (1976)

November 2022. Jake Moffatt, a Canadian who has just lost his grandmother, checks the Air Canada website that same day. The company's chatbot helpfully explains a policy that does not exist: that he can buy the ticket at full price and then claim the bereavement fare within 90 days. Moffatt buys — 1,630.36 Canadian dollars. When he claims, the company replies that the real policy does not apply to travel already taken, and refuses. In February 2024, the British Columbia civil tribunal orders Air Canada to pay him 812 dollars: the company "did not take reasonable care to ensure its chatbot was accurate" [1]. And it leaves a line for the history books: it should be obvious to Air Canada that it is responsible for all the information on its website — it makes no difference whether it comes from a static page or from a chatbot. Along the way, the company had gone so far as to argue that the chatbot was "a separate legal entity," responsible for its own acts. The tribunal did not buy it [1].

That was just a chatbot giving information. What is already here are agents that act: they buy tickets, contract services, answer customers, move money. Every act of an agent is a legal act of whoever activated it — and every error, a three-way dispute: the user, who maintains they never authorized that; the agent's platform, which insists it executed what it was told; and the merchant, who received an apparently valid order.

This chapter deals with delegation to machines that make mistakes. What the model is inside — the data, the weights, the AI-BOM — is dealt with in chapter 79; the action logs of public sector agents, in chapter 107; the identity of who you hire remotely, in chapter 91; the executive impersonation that orders transfers, in chapter 87. The key question: when the purchase goes wrong, how do you prove exactly what you authorized — and exactly what the agent did? The answer: with the mandate and the action log sealed with **t2p** — the evidentiary package of chapter 13 [3] — in a neutral layer that none of the three interested parties controls.

The industry is already building the mandate

Read the landscape through this book's eyes, because it is revealing [2]. Visa, with Intelligent Commerce, issues limited-scope tokenized credentials for agents: a token for the travel agent valid only for one airline and a window of dates; a token with a $200 limit that cannot authorize a $500 purchase. Mastercard announced Agent Pay in April 2025: verified agents transacting with "Agentic Tokens" that bind the payment credential to a specific agent, a scope of merchants and a consent policy. And Google published in September 2025 the open protocol AP2, with more than sixty partners — Mastercard, American Express, PayPal, Coinbase, Salesforce — which defines two pieces with revealing names: the "Intent Mandate" — what the user wants — and the "Cart Mandate" — what the agent proposes to buy — both signed as verifiable credentials [2]. Meanwhile, the agents themselves are already consumer products: OpenAI's Operator, Google's Project Mariner, Anthropic's Computer Use, Microsoft's Copilot agents [2].

Notice the vocabulary: mandates, verifiable credentials, scope, consent. It is the language of age-old civil law, rewritten in cryptography. The payments industry has understood before anyone else that the agents problem is the mandate.

But each solution lives inside its own ecosystem — Visa's watches the Visa payment; Google's, its protocol — and they all watch the moment of payment. The dispute, however, comes afterward and is broader: what you authorized, what the agent saw, what it did, in what order. And the evidence of this cannot live in the house of any of the three interested parties — not on the phone of the user who claims, not on the servers of the platform that defends itself, not in the systems of the merchant that got paid. The neutral layer is missing.

The proof of existence applied

Four gestures turn delegation into an auditable fact [3].

1. The sealed mandate

Before activating the agent, the user — or the company — drafts the mandate: maximum budget, permitted scopes, vetoes, duration. The document is sealed with t2p: SHA-512 fingerprint computed on the device, certain date, anchoring [3]. No one else needs to see it: on the day of the dispute, the mandate existed with this exact content before the agent's first action. The "I never authorized this" and the "yes you did" stop being a contest of memories: there is a text, with a date.

2. The anchored action log

The agent generates its log — what it did, when, on which sites — and this log is sealed periodically in batches: the fingerprint of each stretch, chained to the previous one [3], like the billing records of chapter 34. The platform that offers this as standard turns its agent's black box into an auditable history — and protects itself: against the customer who denies receiving the summary, or the merchant who alleges a phantom order, the anchored log answers.

3. The instant of the offer

Prices and terms change by the second; the agent decides on what it sees now. The certified capture of what the agent had in front of it — the offer page, with full image and source code [3] — freezes the instant of the decision. When the merchant says "this price never existed," or the user says "the offer promised something else," the screen of that moment has a certain date. It is the lesson of Air Canada [1], carried into the world of agents: the information acted upon, preserved.

4. The strong authorization

Below the threshold the user set in the mandate, the agent goes alone; above it — the big purchase, the new contract, the unusual payment — the act requires explicit human approval: t2p's video-selfie signature, with its two random liveness challenges and the device's biometric verification [3]. The approval is bound to the specific act, with a certain date. Neither can the user deny what they approved looking into the camera, nor can anyone pass off as approved what was never submitted.

The limits, stated clearly: the seal does not make the agent smarter — it does not prevent the hallucination or the error: it detects them, dates them and makes them attributable. Final responsibility is allocated by contracts and courts — and Air Canada [1] set the principle: whoever deploys the agent answers for what the agent says and does. And not everything has to be an agent: there are acts — the most personal ones — that deserve human hands from start to finish.

The promise and the peril

The promise: delegation with confidence — the user who activates an agent knowing that the mandate and every action have a certain date [3]; disputes resolved in minutes by comparing mandate, log and capture, instead of years of litigation [1]; platforms competing on auditability; and the mandate infrastructure the industry is already building [2], completed with the neutral layer none of the parties controls.

The perils, faced head-on.

The first is **your whole life in the log**. A personal agent's action log is an intimate diary: what you buy, what you read, what you negotiate. It must stay on the device or in the user's systems — only fingerprints on the chain [3] — with selective disclosure of the disputed stretch, never of the whole life.

The second is **the agent as alibi**. "The agent did it" cannot become the digital version of "it wasn't me." The dated mandate delimits responsibilities, it does not erase them: whoever activates an agent with a broad mandate assumes what fits inside it; whoever deploys it commercially answers to the customer — the Air Canada principle [1]. The trail serves to apportion blame properly, not to make it disappear.

The third is **the agents race**. The merchant has its agents too, and the asymmetry reproduces itself: the giant with a fleet of agents and perfect logs against the consumer with a phone. The neutrality of the proof is the minimum; the protection of the weaker party, as always, is regulatory.

And the fourth is **agentic techno-solutionism**. No seal makes delegating always a good idea. The agent is a tool for delegable tasks; the judgment about what gets delegated — and what does not — remains the most human decision of all.

It's time to act

Let delegation have a memory.

For those who activate agents — individuals and companies: a written, sealed mandate before activating [3]; human-approval thresholds for the big acts; and the log, exported and anchored periodically.

For agent platforms: an auditable, anchored action log as standard [2][3] — your defense in the dispute and your sales argument; and let the user's mandate be a verifiable object, not an accepted checkbox.

For merchants: capture and anchor what you serve to the agents — the offer, the price, the terms of the instant [3]; tomorrow's dispute is won with today's screen.

For payment networks and protocols: your signed mandates [2] are the foundation; open them to neutral anchoring so the proof survives any one of you.

For regulators and legislators: the Air Canada principle [1] — whoever deploys, answers — deserves the rank of law; and every agent acting before consumers should have a dated mandate and an auditable log as a condition of operation.

We delegate because life does not fit in our hands; let delegation have a memory. When your agent buys, contracts and answers for you, may you be able to prove what you ordered it to do — and it, what it did. For delegation with a certain date, **time2act**.

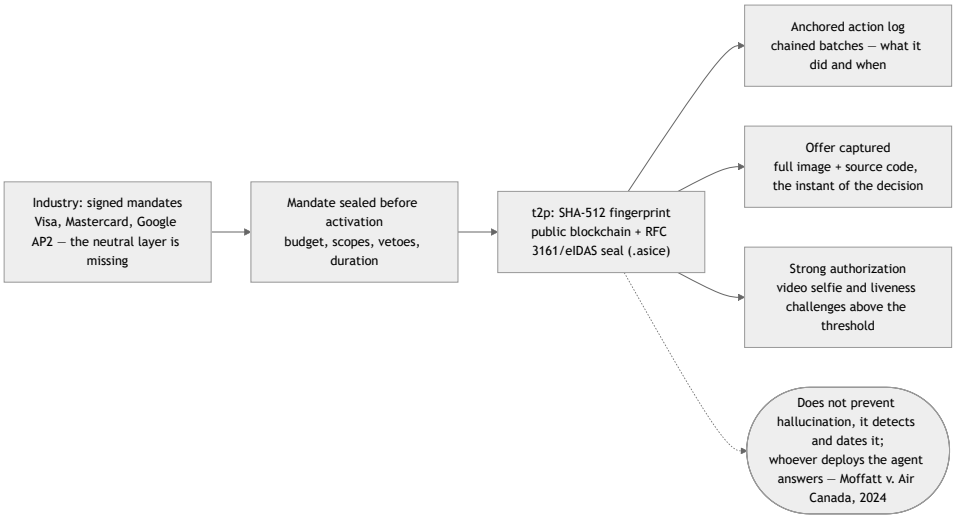


Figure 80.1 — The verifiable mandate for AI agents with proof of existence (sealed with t2p — chapter 13): the mandate sealed before activation (budget, scopes, vetoes), the action log anchored in chained batches, the certified capture of the offer the agent saw (image + source code) and the strong authorization with video selfie above the threshold. Context: the industry is already building signed mandates (Visa Intelligent Commerce, Mastercard Agent Pay, Google AP2 — Intent/Cart Mandate as verifiable credentials); the missing layer is the neutral one between user, platform and merchant. Precedent: Moffatt v. Air Canada (2024) — whoever deploys the agent answers. Sources: British Columbia tribunal; Visa/Mastercard/Google AP2.

---

[1] CBC News — *Air Canada found liable for chatbot's bad advice on bereavement rates* (Moffatt v. Air Canada, B.C. Civil Resolution Tribunal, February 2024: the chatbot invented a bereavement fare policy; CA$812 in compensation on 1,630.36 paid; "Air Canada did not take reasonable care to ensure its chatbot was accurate"; "it is responsible for all the information on its website — whether it comes from a static page or a chatbot"); Forbes — *What Air Canada Lost In 'Remarkable' Lying AI Chatbot Case* (the claim that the chatbot was "a separate legal entity"); McCarthy Tétrault — *Moffatt v. Air Canada: A Misrepresentation by an AI Chatbot*. https://www.cbc.ca/news/canada/british-columbia/air-canada-chatbot-lawsuit-1.7116416 · https://www.forbes.com/sites/marisagarcia/2024/02/19/what-air-canada-lost-in-remarkable-lying-ai-chatbot-case/ · https://www.mccarthy.ca/en/insights/blogs/techlex/moffatt-v-air-canada-misrepresentation-ai-chatbot ↩

[2] Visa — *Intelligent Commerce* (limited-scope tokenized credentials for agents; MCP server; a $200 token cannot authorize $500); Mastercard — *Agent Pay* (announced 29-4-2025: Agentic Tokens binding credential, agent, merchant scope and consent policy; "Agent Pay for Machines," June 2026); Google AP2 (September 2025, 60+ partners: Intent Mandate + Cart Mandate signed as verifiable credentials); Eco/Paz — 2026 agentic commerce guides (runtimes: OpenAI Operator, Google Project Mariner, Anthropic Computer Use, Microsoft Copilot). https://www.visa.com/en-us/solutions/intelligent-commerce · https://www.mastercard.com/us/en/news-and-trends/press/2026/june/mastercard-launches-agent-pay-for-machines.html · https://eco.com/support/en/articles/14839400-what-is-agentic-commerce-the-2026-guide · https://www.paz.ai/glossary/agentic-commerce-protocol-acp ↩

[3] t2p — product manual (certify any file — the mandate and the logs, only the fingerprint travels; web page capture with full image + source code; signature with video selfie, two random liveness challenges and the device's biometric verification; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the models' provenance, in chapter 79; public sector agents, in chapter 107. ↩

Your Face, Your Voice: Image Rights and Digital Replicas

"The serial number of a human specimen is the face, that accidental and unrepeatable combination of features."

— Milan Kundera, Immortality (1990)

In October 2023, **Tom Hanks** had to warn his followers on Instagram about something strange: "There's a video out there promoting some dental plan with an AI version of me. I have nothing to do with it" [1]. Months later, the actress **Scarlett Johansson** explained that she had turned down OpenAI's request to lend her voice to ChatGPT, and that the company had nonetheless launched a voice — "Sky" — that resembled hers so closely that a forensic analysis by Arizona State University found it more similar to Johansson's than **98%** of the other actresses examined; OpenAI withdrew it after receiving letters from her lawyers [2]. In both cases, the person had to come out and deny it **afterward**, with no way to prove, upfront, that they had never given permission. The burden fell on the victim.

This is this chapter's question, distinct from chapter 40's — the authorship of a work — and from chapter 67's — live performance and reproduction rights. Here the **protected work is you**: your face, your voice, your body. And the question is not who did it, but **did you consent, for what use, and can it be proven?** Because generating a digital replica of someone has become trivial, and denying it has not.

A face that is no longer yours

The problem takes three forms. The first is **fraudulent advertising**: famous faces — Hanks, Johansson — but also anonymous people — pasted onto products they never endorsed, in a wave of scams that multiplies because fabricating them takes minutes. The second is the **fear of voice and image professionals**: as we have seen, in 2023 Hollywood's actors went on strike for months, and one of their red lines was precisely **digital replicas**; the agreement that ended the strike requires informed consent to create or reuse a performer's replica, with a "reasonably specific description" of the intended use, and a **new** consent — and new pay — if the replica is carried into another project [3]. The third is the **resurrection of the dead**: voices and faces of the deceased reanimated to sell or entertain without anyone being able to say whether the family authorized it. In all three, the same asymmetry: it is easy to fake and hard to deny, and whoever suffers the harm has to chase every copy one by one.

The law recognizes the right; the receipt is missing

The law has reacted with unusual speed. In Tennessee, the **ELVIS Act** — *Ensuring Likeness, Voice and Image Security* — came into force on July 1, 2024 and brought the **voice** into the right of publicity, even creating liability for whoever distributes an AI tool whose primary purpose is to imitate someone's voice or image [4]. In the United States, the **NO FAKES Act** — which would create a federal right over the digital replica — cleared the Senate Judiciary Committee unanimously on June 18, 2026, though it still has to pass the full Senate and the House of Representatives [5]. And Denmark has taken the most radical step: in June 2025 it introduced a reform of its intellectual property law — still in parliamentary procedure, with entry into force expected in mid-2026 — to give **every person a copyright-style right over their own face, body and voice**, with the right to demand the removal of deepfakes and compensation, protection up to fifty years after death, and exceptions for parody and satire [6]. In Europe, Article 50 of the AI Regulation will require **marking** AI-generated content [see chapter 79]. But all these laws collide, when it matters, with a single question: **did this person consent to this use, and where is the proof?** Today consent is a PDF in a drawer — or it is nothing.

Why the law alone does not save you

It must be said bluntly: a right of publicity is worth only as much as the evidence that sustains it. When an ad with your face appears, you have to prove that you did **not** consent — proving a negative — and the platform must decide within hours whether to take it down. An assignment contract stored on a studio's server can be touched up or backdated; an unscrupulous studio can claim "they already agreed" and produce a piece of paper. Without a consent record that is **verifiable and bears a certain date**, every

dispute becomes your word against theirs, and whoever has more lawyers wins. The right exists; what is missing is the **unalterable receipt**.

The proof of existence applied

This is where **t2p** contributes two complementary layers. The first is controlled by the **person**: with a **proof of life**, they create an *image consent record* that binds a **biometric reference** — the SHA-512 fingerprint of a face or voice sample, which **never leaves the device** — to an explicit, **bounded** grant: which uses they authorize, for whom and until when. All of it is sealed with an **RFC 3161/eIDAS timestamp** and anchored to the **public blockchain** inside an **.asice** container. A legitimate digital replica then carries a **verifiable pointer** to a sealed, current consent; and the **revocation** is also sealed and dated, so that no one can invoke a permission that had already been withdrawn. The second layer is sealed by the **studio or the advertiser** at the moment of signing: the informed consent in the style of the actors' agreement, with the specific description of the use and the performer's proof of life, so that reusing the replica in another project without a new consent becomes **detectable**.

The effect is that the burden of proof switches sides. The Hanks case stops depending on his word: he can show that he **never** issued any sealed grant for that dental ad, and the **absence** of an anchored permission becomes, in itself, proof of the lack of consent. It combines with the provenance of chapters 13 and 15: the C2PA mark or SynthID says "this is synthetic"; the consent record adds "and no one authorized this face."

Now the limit, in all its rawness. "Garbage in, garbage on-chain": the seal proves that consent was given — or not — by whom and when; **it does not detect the deepfake** on its own (that depends on the provenance mark and the detectors), **it does not prevent its creation**, and a consent sealed under pressure is still a vitiated consent, now immutable besides. It proves the **act** of consenting, not its **justice**. What it does is take the disputes out of the terrain of word against word and into that of "show the sealed grant."

The promise and the peril

The promise is a world where your face and your voice are **provably** yours again: performers license their replicas with a dated, bounded, revocable consent; platforms and courts have a record to distinguish authorized use from violation; and the dead are protected by grants sealed by their heirs. In an age when any face can be cloned in minutes, making consent immutable is returning control to the person.

The peril calls for nuances. First, a consent **extracted by coercion** and sealed does not stop being unjust because it is immutable. Second, building a registry of biometric references is sensitive: hence the golden rule — the face **does not leave the device**, only a fingerprint is sealed — must replace any central database of faces, which would be a dream for surveillance. Third, studios must not end up holding everyone's image files. And fourth, the absence of consent must not be confused with legitimate uses — parody, satire, criticism, reporting — which the law already protects. The proof of existence is a layer of control over one's own image, not an automatic permission or a falsehood detector.

Moreover, delegating biometric management to decentralized networks does not eliminate the risk that effective control concentrates in the hands of a few technical actors. The infrastructure needed to create and validate fingerprints requires powerful servers, access to massive databases and a complex software layer that only a limited number of companies can develop and maintain. As Marc Goodman points out, Facebook — and not a state program — is "the most extensive repository of biometric data on Earth, far surpassing that of any government" [11]. This shows how a private platform can accumulate biographical information on a global scale, becoming a privileged entry point for the capture and exploitation of the fingerprints. Therefore, even with protocols guaranteeing that the face does not leave the device, independent governance and audit mechanisms must be established to prevent this technical concentration from translating into real power over people's identity.

As we have seen, the example of **PalmerAI** in the American Samoa Democratic caucus illustrates how the question is no longer "is the video real?", but "what did the person say and what does the replica say?" [12]. The advance of an avatar that reproduced Jason Palmer's voice and appearance, capable of winning an election 51-40, highlights a gap: the voter needs to know whether the message comes from the candidate or from a digital entity. A neutral blockchain anchor can register each statement with an immutable proof of existence, linked to the original consent certificate and to the holder's private key. Thus any observer

could trace the origin of the speech and validate it without depending on centralized platforms. However, the same mechanism can create a new risk: if the keys are compromised or wielded by a power group, the certifying chain could be used to legitimize manipulated replicas, becoming a "trust layer" that, badly governed, reinforces the concentration of control over political and social discourse.

It's time to act

For **individuals and performers**: seal a bounded, revocable image consent record, keeping the biometric reference private on the device. For **studios, advertisers and AI tool providers**: seal the informed consent at signing, with the specific scope of the use, and **verify a current grant** before generating any replica. For the **platforms** where the ads run: require a verifiable consent pointer for any synthetic endorsement, and treat its absence as a takedown signal. And for **legislators and courts** (the ELVIS Act, NO FAKES, Denmark, the European Regulation): recognize the sealed consent record as evidence with a certain date of the consent — or of its lack.

Tom Hanks had to come out and say that the face selling dental plans was not his. Giving every person a sealed proof of what their image authorizes, for whom and until when — never forgetting that the seal proves the consent, not its justice — is, on this front where our faces can be stolen, **time2act**.

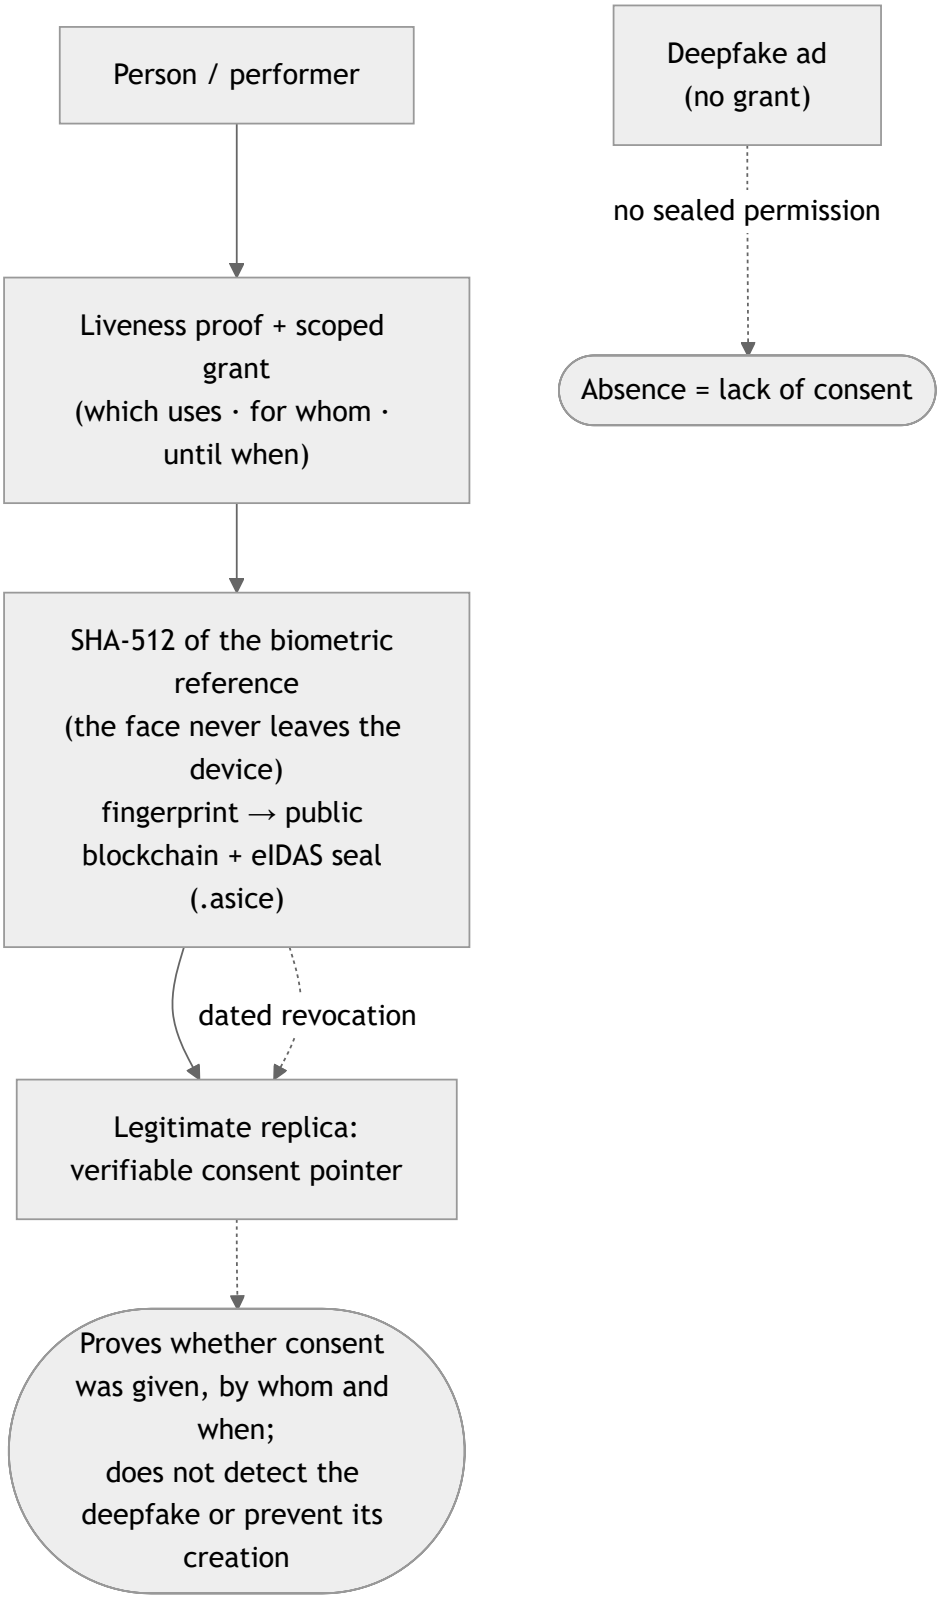


Figure 81.1 — Two paths. Above: the person creates, with proof of life, a bounded consent grant (which uses, for whom, until when) bound to a biometric reference (SHA-512 fingerprint; the face does not leave the device), anchored to the public blockchain, with a complementary eIDAS seal (.asice); a legitimate replica carries a verifiable pointer, and the revocation is dated. Below: a deepfake ad with no sealed grant — the absence itself proves the lack of consent. Limit: the seal proves whether there was consent, from whom and when; it does not detect the deepfake or prevent its creation, and it does not judge whether the consent was free. Sources: the Tom Hanks case; the Johansson–OpenAI case; the SAG-AFTRA agreement; the ELVIS Act; the Danish reform.

[1] Voicebot.ai — *Tom Hanks Warns of Deepfake Hanks Dental Ad* (October 2023; "I have nothing to do with it"). https://voicebot.ai/2023/10/02/tom-hanks-warns-of-deepfake-hanks-dental-ad/ ↩

[2] NPR — *OpenAI pulls AI voice that was compared to Scarlett Johansson in the movie 'Her'* (May 2024; forensic analysis by Arizona State University: more similar than 98% of actresses). https://www.npr.org/2024/05/20/1252495087/openai-pulls-ai-voice-that-was-compared-to-scarlett-johansson-in-the-movie-her ↩

[3] SAG-AFTRA / Perkins Coie — *How the 2023 SAG-AFTRA contract addresses generative AI and digital replicas* (informed consent; specific description of the use; new consent for reuse). https://perkinscoie.com/insights/blog/generative-ai-movies-and-tv-how-2023-sag-aftra-and-wga-contracts-address-generative ↩

[4] Davis Wright Tremaine — *Tennessee Passes ELVIS Act (Ensuring Likeness, Voice and Image Security Act)* (in force July 1, 2024; the voice within the right of publicity; AI tool liability). https://www.dwt.com/blogs/artificial-intelligence-law-advisor/2024/04/tennessee-elvis-act-ai-voice-replica ↩

[5] United States Senate Judiciary Committee — *NO FAKES Act (Nurture Originals, Foster Art, and Keep Entertainment Safe)* (a federal right over the digital replica; revised 2026 version, approved by the Senate Judiciary Committee on June 18, 2026). https://www.judiciary.senate.gov/no-fakes-act ↩

[6] World Economic Forum — *Deepfake legislation: Denmark takes action* (introduced June 2025; a copyright-style right over face, body and voice; removal and compensation; parody exceptions; in parliamentary procedure, entry into force expected mid-2026). https://www.weforum.org/stories/2025/07/deepfake-legislation-denmark-digital-id/ ↩

[7] Variety — *Scarlett Johansson Reacts to OpenAI Sky Voice: 'Shocked' and 'Angered'* (May 2024). https://variety.com/2024/digital/news/scarlett-johansson-responds-shocked-angered-openai-chatgpt-her-1236011135/

[8] SAG-AFTRA — *Digital Replicas 101* (typology and the performer's consent rights). https://www.sagaftra.org/sites/default/files/sa_documents/DigitalReplicas.pdf

[9] Wikipedia — *ELVIS Act* (Ensuring Likeness, Voice and Image Security Act of 2024; reform of the Personal Rights Protection Act of 1984). https://en.wikipedia.org/wiki/ELVIS_Act

[10] EU Artificial Intelligence Act — *Article 50: transparency obligations* (marking of AI-generated or manipulated content). https://artificialintelligenceact.eu/article/50/

[11] Marc Goodman — *Delitos del futuro* (Ariel, 2015). ↩

[12] Wikipedia, *Jason Palmer (politician)*; TIME, *Albania's AI minister Diella* (2025). https://time.com/7324934/albania-ai-minister-diella/ ↩

Sharenting: The Image of Those Who Cannot Consent

"There are no children — there are people; but with a different scale of concepts, a different store of experience, different impulses, a different play of feelings."

— Janusz Korczak, When I Am Little Again (1925)

The first day of school, with the new backpack and the half-formed smile. The birthday party, the balloons, the cake. The afternoon at the beach. And the automatic gesture every time: post it. No malice — love, pride and the century's habit. The datum that frames it [1]: a child appears on average in 1,300 photographs published online before turning 13 — a public album they did not choose, indexed, replicable and permanent. And the risks the French legislator listed while processing its law are not hypothetical: identity theft, blackmail, cyberbullying — and the worst ones [1].

The nuance that makes the whole chapter: this is not about blaming parents. It is about a legal question no one had asked until recently — who consents to the image of someone who cannot yet consent, and where is it written down? Who can prove that a post had the permission of both parents, or that one refused? Where is the authorization the school says it had?

This chapter deals with consent over minors' images — and how to attest it. Adults' image rights — the protected work is you — are in chapter 81; the dissemination of intimate images, in chapter 83; hate and harassment against the young, in chapter 84; co-parenting — where the minor is never the object of the record — in chapter 53; school, in chapter 19. The key question: how do you protect the right of someone who cannot yet exercise it? The answer: by documenting the decisions of the adults who hold it in custody — the authorization, the refusal, the takedown request — with a certain date: **t2p**, the usual mechanics, that of chapter 13 [3]. Consents are documented; never the minor — chapter 53's line is, here, sacred.

France writes what no one had written

On February 19, 2024, the French Parliament passed — unanimously — Law 2024-120, guaranteeing respect for children's image rights [1]. The reform touches the heart of family law: Article 371-1 of the Civil Code now says, expressly, that parental authority includes protecting the minor's private life and their right to their own image; that the parents protect it jointly; and that they must involve the child according to their age and maturity. And the most forceful piece: if the parents' own dissemination of their child's image gravely harms the child's dignity or moral integrity, the judge can delegate the exercise of the minor's image right to a third party — taking it away from the parents [1].

In Spain, the framework protects the minor intensely — the organic law protecting honor, privacy and one's own image, the law on the legal protection of minors, the oversight of the public prosecutor [2] — and the courts already require the agreement of both parents to publish images of the children: one parent's refusal should stop the dissemination [2].

And here, this book's diagnosis: the legal consensus exists; its proof does not. The whole edifice — the joint consent, one parent's refusal, the school authorization, the takedown request — runs on acts no one can attest afterward: the separated father's "I never authorized this," the "the school had permission" that cannot be found, the "I asked them to delete it a year ago" with no date. Statements that weigh little when the conflict comes — because they have neither date nor form.

The proof of existence applied

The decisions of the adults are documented — never the minor [3].

1. The authorization that says what

The school or after-school authorization — often a generic checkbox — turned into a precise, sealed document: which uses (the school website? social networks? the families' group?), which channels, for how long; signed with a certain date [3]. The school that has it anchored knows what it may do; the family knows what it has allowed; and "we had permission" becomes a consultable fact, not a start-of-term argument.

2. The disagreement on record

When one parent does not want images of the child published — the classic case in a separation — their refusal, communicated and sealed with a certain date [3]: to the other parent, to the school, to the club. The courts already require the agreement of both [2]; what was missing was for the refusal to have a date. The logic is chapter 53's — documenting the adults' agreements and disagreements, never the minor — applied to the image.

3. The dated takedown

The request to delete a photo — to the other parent, to the enthusiastic grandfather, to the platform — sent and captured with a certain date; and the silence or the refusal, dated too [3]. If the conflict escalates — up to the judge the French law already provides for [1] — the whole chronology exists: it was requested on such a day; it was not honored; this is what was published.

4. The fingerprint for the future

The published photo — or the whole album — registered by fingerprint: the SHA-512 on the family's device, none of the content anywhere [3]. It might seem paradoxical — documenting what we wish did not exist — but it is the evidentiary gift to the grown child: the day they want to exercise their rights — removal, oblivion, or simply knowing — they will be able to attest what was published about them and when, without depending on anyone's memory or on any search engine's cache. Their digital childhood, auditable by them.

The limits, before anyone gets their hopes up: the seal attests to consents, refusals and dates — it does not decide what the best interests of the child are: that is done by parents who get along and, when they do not, by the judge [1][2]. None of this is an invitation to document the child — chapter 53's line rules: what gets recorded are the adults' decisions. And the best protection remains the oldest: post less, think first, and ask the child when they are old enough to answer — the French law says it with a beautiful word: involve them [1].

The promise and the peril

The promise: schools with precise, consultable authorizations [3]; separated parents who argue the substance — what is good for the child — and not the facts; takedown requests that do not get lost in conversations; and a generation that, on turning 18, will be able to look at their digital childhood with their rights in hand [1] — because someone took the care to date it.

Now, the perils.

The first is **the dossier against the ex-partner**. The documented disagreement can become ammunition in the parental conflict. The counterweight is chapter 53's: the trail is to protect the minor, not to win the war — the bilateral design and the child's best interests as the compass; and judges quickly spot who documents for the child and who against the other parent.

The second is **the bureaucratization of tenderness**. If posting the birthday photo requires a procedure, family life becomes impossible. The answer is proportionality: the precise consent is for institutions — schools, clubs, media — and for disagreements; between adults who get along, respect is enough.

The third is **the false seal of legitimacy**. A dated authorization does not make a harmful post good. Parental consent has limits — the child's dignity draws them, and the French law has written it down [1]: there are disseminations no parental signature covers.

And the fourth, the most intimate: seeking in a tool the excuse not to think. No fingerprint replaces judgment. The question before posting is not "do I have permission?", but "will they be glad, fifteen years from now, that this exists?" The proof of existence orders the consents; prudence makes them unnecessary.

It's time to act

Let every image of a child have behind it a decision that can be shown.

For mothers and fathers: think before posting — and when the child is old enough, ask them [1]; if you disagree between yourselves, let it be on record with a date [3]: the child deserves that their image not be a battlefield.

For schools, clubs and organizations: precise, sealed, renewed authorizations [3] — the generic start-of-term checkbox protects neither the families nor the institution.

For the platforms: takedown requests for minors' images with a dated, traceable response — silence is a decision too, and it should be on record.

For legislators: France has shown the way [1] — the minor's image right as a joint duty of the parents, with the judge as the last guarantee; let the transposition of this idea come with attestable consents and refusals.

For the children, when you are grown: your digital childhood is yours; claim it — and may the dated fingerprints [3] spare you having to prove the obvious.

The 1,300 photos [1] are already out there; the next ones are not yet. Let every image of a child have behind it a consent that can be shown — and an adult who stopped to think. For those who cannot yet say no, **time2act**.

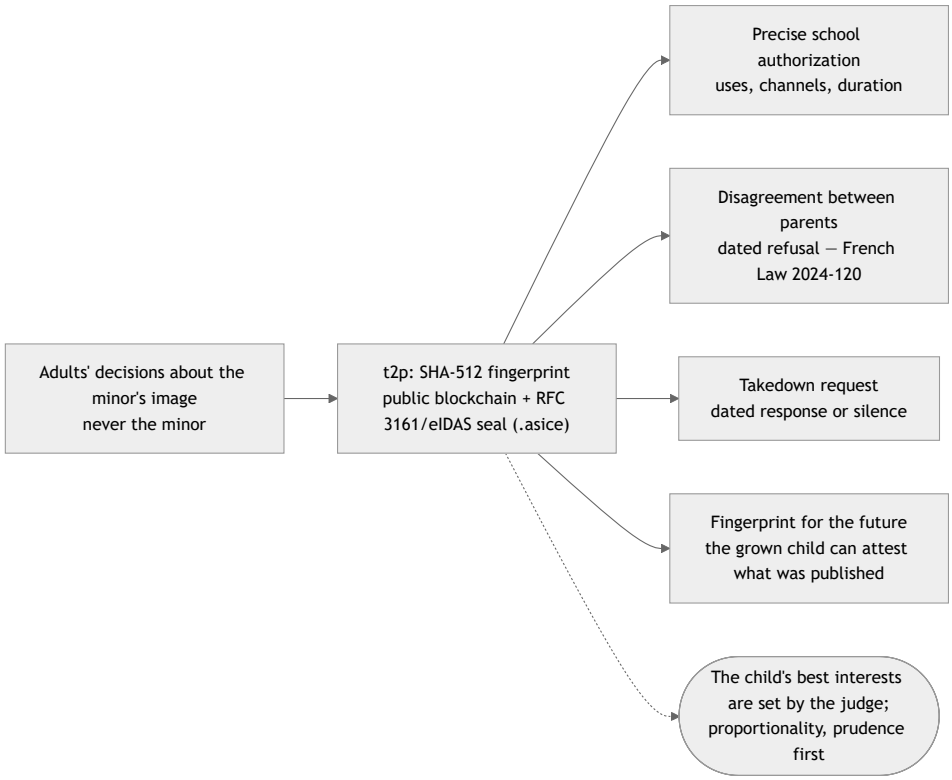


Figure 82.1 — The minor's image with proof of existence (sealed with t2p — chapter 13) — the decisions of the ADULTS are documented, never the minor: the precise, sealed school authorization (uses, channels, duration), the dated disagreement between parents, the takedown request with dated response or silence, and the fingerprint of the post for the future (the grown child will be able to attest what was published about them, without any content traveling). Context: a child appears in ~1,300 photos online before age 13; French Law 2024-120 (the image right within parental authority, exercised jointly; judicial delegation in grave cases). Limits: the child's best interests are set by the judge; proportionality; prudence first. Sources: Légifrance; DAJ.

---

[1] Légifrance — *Law No. 2024-120, of February 19, 2024, guaranteeing respect for children's image rights* (amendment of Art. 371-1 of the Civil Code: parental authority includes protecting the minor's private life and image right; protection exercised jointly by both parents, involving the child according to age and maturity; judicial delegation of the exercise of the image right in case of grave harm to dignity or moral integrity); DAJ/economie.gouv.fr — *Respect du droit à l'image des enfants: les apports de la loi* (context: a child appears on average in 1,300 photographs online before age 13; risks: identity theft, blackmail, cyberbullying); Actu-Juridique — *Les 5 apports de la loi.* https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000049163317 · https://www.economie.gouv.fr/daj/lettre-de-la-daj-respect-du-droit-limage-des-enfants-les-apports-de-la-loi-du-19-fevrier-2024 · https://www.actu-juridique.fr/civil/personnes-famille/le-respect-du-droit-a-limage-des-enfants-et-les-5-apports-de-la-loi-du-19-fevrier-2024/ ↩

[2] The Spanish framework (LO 1/1982 on honor, privacy and one's own image; LO 1/1996 on the legal protection of minors; intervention of the public prosecutor; case law requiring the agreement of both parents to publish images of the children). See chapter 81 for general image rights and chapter 53 for co-parenting records. ↩

[3] t2p — product manual (document scanner; certified screenshot and email; certify any file — only the fingerprint travels, the content never leaves the device; signature with video selfie; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; co-parenting, in chapter 53; intimate dissemination, in chapter 83. ↩

Non-Consensual Intimate Images: The Harm, the Proof and the Takedown

"To girls everywhere, I am with you. On nights when you feel alone, I am with you. When people doubt you or dismiss you, I am with you."

— Chanel Miller, victim impact statement, People v. Turner (2016)

In September 2023, in **Almendralejo**, a town in Extremadura, dozens of teenage girls discovered that nude images of them were circulating on WhatsApp — images that had never existed. Someone had taken clothed photos from their Instagram accounts and, with an AI application, turned them into fake nudes. The victims were between twelve and fourteen years old. "Many girls were completely terrified and had tremendous anxiety attacks because they were suffering it in silence," one mother explained [1]. A juvenile court in Badajoz convicted fifteen boys — thirteen to fifteen years old — of twenty offenses of creating child abuse material and twenty more against the victims' moral integrity [2]. The harm was already done: the images were fake, but the humiliation and the fear were real.

This chapter deals with one of the cruelest faces of deepfakes, distinct from chapter 81's — the commercial use of one's image. Here we speak not of money or brand, but of **dignity**. And one figure must be said because it frames everything: more than **90% of the deepfake content** circulating on the network is pornographic, and nearly all of it has a **woman** as its victim [6]. This chapter's question is not how to prevent it from existing — that cannot be done entirely — but how to give the victim **proof of the harm and a provable takedown** without forcing her to relive it.

A harm easy to inflict, hard to repair

The Almendralejo case is not isolated. In January 2024, fake sexual images of **Taylor Swift** flooded the network X: a single post was viewed more than **47 million times** before it was taken down, and the platform blocked searches of her name for two days [3]. If this happens to one of the most powerful women on the planet, imagine an anonymous teenager. The asymmetry is brutal: fabricating the image takes minutes and is free; removing it forces the victim to chase every copy on every platform, prove that it is her, and often **look again** at the thing that wounds her in order to submit it as evidence.

A defense already exists that, not by chance, uses exactly this book's principle. **StopNCII.org** lets a victim generate on her device a **fingerprint** (a hash) of the intimate image; **only the fingerprint — never the image — leaves the phone**, and it is shared with Meta, TikTok, Reddit or Google so they block matching uploads. Since 2021 it has protected more than **500,000 images** and prevented more than **10,000 attempted** publications, with a success rate above 90% [7]. It is the proof that the victim can be defended **without centralizing** a single image.

The law is in place; the unalterable proof is missing

The legal response has been fast and forceful. In the United States, the **TAKE IT DOWN Act**, signed on May 19, 2025, criminalizes the publication of non-consensual intimate images — deepfakes included — and obliges platforms to remove them within a maximum of **48 hours** of notice; since **May 19, 2026** the FTC has been enforcing compliance on the platforms [4]. In Europe, **Directive (EU) 2024/1385** on violence against women, in force since June 2024 and to be transposed by June 2027, criminalizes the non-consensual production and dissemination of manipulated intimate material and obliges states to provide **takedown orders and rapid preservation of evidence** [5]. The law, then, already demands two things: remove fast and **preserve the evidence**. And here the usual wall returns: a screenshot is manipulable, the notice to the platform may leave no dated trace, and the victim ends up having to keep custody, herself, of the material that wounds her.

Why the report alone does not save

It must be said bluntly, because here it is especially delicate. For a court to convict — as in Almendralejo — and for a takedown to be enforceable, **admissible evidence** is needed: that this image existed, where, when, and that its removal was requested. Today this evidence is fragile: a capture without a seal does not

prove the date; proving that the platform received the notice and failed to act within 48 hours depends on emails and records that can be disputed; and preserving the evidence means, too often, that the victim has to store and show again her own image. What is missing is a way to **fix the proof of the harm and of the notice with a certain date, without reliving it or storing it**.

The proof of existence applied

This is where **t2p** extends StopNCII's logic with the missing piece. When a victim, a support organization or a platform detects a non-consensual image, the **SHA-512 fingerprint** of the content is computed on the device and sealed, together with **where and when** it was found (the URL, the time) and the **takedown request**, with an **RFC 3161/eIDAS timestamp** anchored to the **public blockchain** inside an **.asice** container. What results is an **unalterable, dated chain of custody** — useful for the prosecution, as in the Almendralejo case, and for proving that the platform was notified at such a moment and had 48 hours to act — **without the victim having to store or share the image again**: only its fingerprint is sealed. The certain date proves that *that* content existed and was flagged as abuse at *that* instant, and the accused can no longer claim the evidence was fabricated afterward.

There is a second application, for the inverse face of the problem: **sextortion** and the false accusation. When someone threatens a victim with disseminating an image that looks real, **provenance** — the absence of any sealed origin, combined with the C2PA mark of chapters 13 and 15 — helps establish that the image is **synthetic**, and strips it of its blackmail power.

Now the limit, in all its rawness, because deceiving oneself here would do harm. "Garbage in, garbage on-chain": the seal **does not remove the image**, **does not prevent its creation**, does not by itself **detect** that it is a deepfake, and cannot make what has been seen be unseen. What it does is provide **proof of the harm and a provable takedown**, not prevention. And it imposes an absolute red line: it can **never** become a registry of intimate images; only fingerprints computed on the device, as in StopNCII. Real protection remains the sum of the fingerprint, the law and the platform's duty; the seal adds the unalterable proof.

The promise and the peril

The promise is to give the victim back a little power: a **dated, admissible proof** of the harm and of the notice, and a takedown that is **enforceable and verifiable** across all platforms, without having to surrender a single image or relive it each time. The aggressor and the slow platform lose the option of denial; the girl from Almendralejo would have, tomorrow, a proof no lawyer could dispute.

The peril demands the book's strictest nuances. First: **no centralized images**; only fingerprints on the device. Second: the reporting process must not become a new source of trauma; the tool must accompany, not force one to look. Third: the governance of the fingerprint bank, which must be audited and abuse-proof. And fourth: always remember that the seal proves **the act of flagging** a piece of content as non-consensual, with date and author; determining who created it belongs to the investigation. The proof of existence is a tool of dignity and justice, not a technological shield that erases the harm.

Moreover, the immutability inherent to blockchain-based records can become an invisible chain of persecution for victims seeking to rebuild their lives after the trauma. Once the fingerprint of a non-consensual image is recorded, it is inscribed forever in the public ledger of evidence; the fact that it cannot be deleted also means there is no natural mechanism to withdraw it when circumstances change or when the victim decides to leave that chapter behind. As Edward Snowden warns, the existence of a "past I entirely regretted, but could barely remember" can remain "so inextricably and technologically bound" to the affected person's identity, increasing the risk of long-term revictimization [11]. That is why any digital sealing architecture must include a layer of expiry or revocability controlled by an independent authority, allowing the fingerprint to be erased or invalidated when the proof is no longer needed or becomes counterproductive. This safeguard would balance the need for legal certainty with the right to be forgotten and to personal rehabilitation.

As we have seen, the magnitude of the problem already exceeds the private sphere and reaches the heights of politics: the American Sunlight Project found that, in December 2024, close to one in six members of the U.S. Congress are victims of AI-generated non-consensual intimate images, a figure that reveals a systematic pattern of abuse and vulnerability at the institutional scale [12]. This reality obliges platforms to adopt "Take It Down" mechanisms capable of responding in real time, but without a proof of existence

anchored to the blockchain the takedown remains up in the air: there is no immutable record proving that the report was made and the content removed. The neutral anchoring of the evidence, independent of the creator and of third parties, turns every removal request into a verifiable lead for the court, sparing victims the burden of proving the harm and guaranteeing that the authorities can act with the necessary legal certainty.

It's time to act

For victims and support organizations (StopNCII, helplines): seal the fingerprint and the circumstances of the harm as evidence, without storing the image. For platforms: maintain the fingerprint bank and issue sealed receipts for every notice and takedown, proving compliance with the 48 hours. For police and courts: treat the timestamped, sealed fingerprint as a chain of custody, as was done in the Almendralejo case. For legislators (the TAKE IT DOWN Act, Directive 2024/1385): recognize the sealed proof as valid evidence and guarantee its rapid preservation. And for the providers of AI apps that "undress" photos: take responsibility for the tool, as the new image laws already demand.

In Almendralejo, girls suffered in silence a harm made with one click. Giving every victim a sealed proof of the harm and a takedown that can be enforced — without ever centralizing her image or forcing her to relive it — is, on this front that touches the most intimate dignity, time2act.

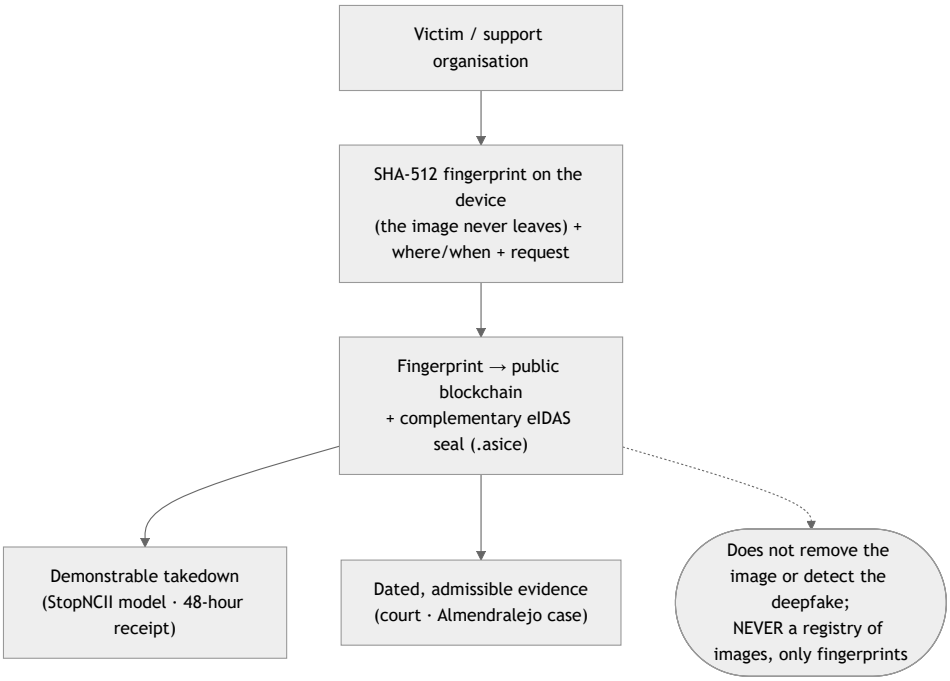


Figure 83.1 — The victim or the support organization computes on the device the SHA-512 fingerprint of the content (the image does not leave the phone) and adds where and when it was found; the whole's fingerprint is anchored to the public blockchain, with a complementary eIDAS seal (.asice). Two outcomes: on the left, the provable takedown across all platforms (the StopNCII model, with a sealed 48-hour receipt); on the right, dated, admissible evidence for the court (as in Almendralejo). Limit: the seal does not remove the image, does not prevent its creation and does not detect the deepfake; never a registry of images, only fingerprints. Sources: the Almendralejo case; the Taylor Swift case; the TAKE IT DOWN Act; Directive (EU) 2024/1385; StopNCII.org.

[1] Euronews — *The Almendralejo case: when AI deepfakes are used to undress teenagers* (September 2023; a mother's testimony; victims aged 12-14). https://www.euronews.com/next/2023/11/15/the-case-of-almendralejo-when-deepfakes-are-used-to-undress-teenagers-euronews-tech-talks- ↩

[2] Scottish Legal News / Reuters — *Spain: court punishes schoolboys for spreading AI deepfakes of girls* (July 2024; 15 minors convicted; 20 offenses of child abuse material and 20 against moral integrity). https://www.scottishlegal.com/articles/spain-court-punishes-schoolboys-for-spreading-ai-deepfakes-of-girls ↩

[3] Wikipedia — *Taylor Swift deepfake pornography controversy* (January 2024; one post viewed more than 47 million times; X blocks searches of her name). https://en.wikipedia.org/wiki/Taylor_Swift_deepfake_pornography_controversy ↩

[4] Congress.gov / The White House — *TAKE IT DOWN Act* (S.146; signed May 19, 2025; takedown within 48 h; platform compliance enforced by the FTC since May 19, 2026; first criminal conviction in April 2026). https://www.congress.gov/bill/119th-congress/senate-bill/146 ↩

[5] Council of Europe — *Directive (EU) 2024/1385 of 14 May 2024 on combating violence against women and domestic violence* (in force June 13, 2024; criminalizes manipulated intimate material; takedown orders and preservation of evidence). https://www.coe.int/en/web/cyberviolence/-/european-parliament-and-council-directive-eu-2024-1385-of-14-may-2024-on-combating-violence-against-women-and-domestic-violence ↩

[6] Euronews — *Nude deepfakes: is the EU doing enough to tackle the issue?* (2024; more than 90% of deepfake content is pornographic, almost all targeting women). https://www.euronews.com/next/2024/03/17/nude-deepfakes-is-the-eu-doing-enough-to-tackle-the-issue ↩

[7] StopNCII.org — *How StopNCII.org works* (fingerprint generated on the device; the image never leaves; more than 500,000 images protected; partners Meta, TikTok, Reddit, Google). https://stopncii.org/how-it-works/ ↩

[8] Skadden — *'Take It Down Act' requires online platforms to remove unauthorized intimate images and deepfakes when notified* (analysis of the notice-and-takedown mechanism). https://www.skadden.com/insights/publications/2025/06/take-it-down-act

[9] RAINN — *TAKE IT DOWN Act* (the perspective of victims of sexual violence). https://rainn.org/federal-legislation/take-it-down-act/

[10] EU Artificial Intelligence Act — *Article 50: transparency obligations* (marking of AI-generated or manipulated content). https://artificialintelligenceact.eu/article/50/

[11] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

[12] The Markup (2024). https://themarkup.org/artificial-intelligence/2024/12/11/1-in-6-congresswomen-targeted-by-ai-generated-sexually-explicit-deepfakes ↩

Hate on the Networks: The Harm, the Proof and Dignity

"If you put shame in a Petri dish, it needs three things to grow exponentially: secrecy, silence and judgment. If you douse it with empathy, it can't survive."

— Brené Brown, "Listening to Shame," TED (2012)

In 2013, in Novara, Italy, Carolina Picchio, a 14-year-old girl, watched a video of herself — apparently drunk at a party — spread across Facebook. The clip unleashed a wave of insults that filled her inbox with some 2,600 messages of mockery, threats and defamation, sent by her ex-boyfriend and the group of friends who backed him. The psychological pressure was so intense that Carolina threw herself from the window of her home. The case became the legislative impulse behind Law 71/2017, the first specific European statute against cyberbullying, championed by Senator Elena Ferrara and dedicated to the victim [1]. This episode shows how digital words, even though they can be deleted or edited in an instant, leave a real and devastating mark on a young person's life.

This chapter focuses exclusively on hate, harassment and humiliation on social networks — not on sex or the dissemination of intimate content (chapters 83 and 81); the general framework of the crisis of truth is in chapter 10, while bot swarms and the civic personhood proof are explained in chapter 108. The key question: how can we offer victims — often minors — tangible **proof** of the harm and a provable takedown without forcing them to keep custody of the content that wounded them? The answer lies in **t2p**, the blockchain-based proof of existence that seals on the device the SHA-512 fingerprint of the hateful content, together with the URL and the detection timestamp, and registers the takedown request. This dated chain of custody is unalterable and admissible in court, all without storing or re-exposing the offensive material [15]. The technology supplies the missing piece: reliable proof of the harm that does not depend on centralized platforms and that protects the victim's dignity while remaining totally neutral.

A harm easy to inflict, hard to repair

Hate on the networks spreads as easily as a text message: a few clicks and an internet connection are enough to set off an explosion of insults, rumors or threats. Repair, by contrast, is an endless walk backward. Every copy of the offensive content replicates across thousands of devices, is stored on servers that often provide no trace of origin and, even when the victim manages to make the original post disappear, the digital traces persist in collective memories and in algorithms that recommend similar content. The burden falls entirely on the person who suffered the aggression: each fragment must be located, shown to be an extension of the harm and, often, the material must be looked at again in order to identify oneself as the target of the hate — a process that reopens the psychological wounds and consumes money and time.

This asymmetry is especially stark among the young. A Pew Research Center study shows that 46% of American teenagers (13-17) have suffered at least one of six forms of cyberbullying, with 32% enduring name-calling and 10% physical threats [2]. The World Health Organization, through the HBSC project, confirms the trend in Europe: 16% (roughly one in six) of young people aged 11, 13 and 15 have been victims of cyberbullying, and the incidence rose between 2018 and 2022 [3]. In Spain, the joint report of the ANAR and Mutua Madrileña foundations indicates that 12.3% of students say they have experienced bullying or cyberbullying, up from 9.4% the year before [4]. In these cases, 91.7% of the harassers are classmates from the same school and 14.2% already use artificial intelligence to create fake videos or impersonations. The impact translates into a palpable reality: the ADL reports that 50% of teenagers have suffered online harassment in the past twelve months, and 36% say the experience spilled over into in-person harassment [5].

These numbers are not mere statistics; they are testimonies of a crisis of trust that extends beyond cyberspace. When verbal violence becomes a form of social control, the mental well-being of the young is compromised and their educational and professional opportunities erode. The magnitude of the problem is plain: this is not a marginal or occasional phenomenon, but an ambient threat that affects collective health and the social fabric. Every time a victim has to chase a digital trail to prove their harm, society abdicates its own responsibility to protect the most vulnerable and perpetuates a cycle of silence and impunity. The solution cannot be limited to punishing the aggressor; what is needed is reliable proof of the harm that is independent of the platforms, neutral and admissible in court — exactly what t2p offers.

The invisible wound and the honest debate

On September 30, 2022, the forensic ruling of Andrew Walker broke a silence that had seemed impregnable: the death of Molly Russell, a 14-year-old British girl, was linked directly to the harmful effects of online content. Of the 16,300 posts the girl had saved on Instagram in the six months before her suicide, 2,100 showed themes of self-harm, depression and incitement to self-destruction [6]. This first ruling of its kind in the world showed that this was no mere tragic coincidence, but content that had eroded Molly's mental health to the lethal point.

The current academic debate revolves around two opposing perspectives. Jonathan Haidt maintains that the massive proliferation of mobile phones and digital platforms has caused a "great rewiring" of the adolescent brain, fueling a mental health epidemic among the younger generations [7]. In opposition, Candice Odgers rejects this direct correlation and argues there is no empirical evidence that the networks "rewire" the brain or that they are the sole cause of the observed disorders; she points instead to social, economic and family factors as complex contributors [7].

Online hate manifests as an invisible wound that, however ephemeral or digitally manipulated, leaves deep scars. The DSA guidelines (2025) impose age verification and pursue harmful content, while the Spanish draft organic law (2024) creates the online restraining order as an aggravating measure [10][11]. Yet these rules depend on reliable proof of the harm: without immutable evidence, judges have only fragile proof at their disposal.

Here t2p supplies the missing piece: immutable proof of the harm — sealed on the device, its fingerprint anchored to the public blockchain inside an .asice container [15] — that allows the authorities to act with certainty and restore, at least in part, the dignity of young victims of digital hate.

Hate no one wrote: deepfakes and swarms

In January 2024, at Pikesville High School (Maryland), a simple audio clip became the most corrosive weapon of a synthetic hate. Dazhon Darien, the school's athletic director, used an artificial intelligence tool to create an audio deepfake that made principal Eric Eiswert appear to utter racist and antisemitic remarks. The audio, spread across social networks within minutes, unleashed criticism and threats against the principal and made him the object of defamation and reprisal. The police identified the content as fake and arrested Darien in April 2024, but the immediate damage could not be reversed: the comments and screenshots had already circulated, feeding a narrative of hate that no one had written [8]. The case shows how synthesis technology can fabricate "words no one ever said" and turn hate into a manipulable product, completely disconnected from the person impersonated.

The tool's other face is the massive amplification of already existing hate speech. A study published in *PNAS* in 2025, focused on the platform X under Musk's leadership, revealed that the weekly rate of hateful content had risen roughly 50% after the acquisition, and that the "likes" associated with this kind of post had increased 70% [9]. It is not always a real community: often what looks like a crowd is the maneuver of a single troll or a coordinated group of fake accounts, capable of multiplying the visibility of the hateful message with few interactions. This dynamic of weak authenticity feeds the crisis of trust described in chapter 108, where the civic personhood proof attempts to reveal who is really behind a digital profile.

In both scenarios, the harm is tangible and real, even if the content is ephemeral or manipulable. Current laws — such as Art. 510 of the Spanish Criminal Code or the United Kingdom's Online Safety Act — punish incitement to hatred and harassment, but they depend on robust proof that is often missing: modifiable screenshots, notices with no time trace and content the aggressor can delete. Here enters *t2p*, applying the forensic recipe of chapters 13/15: a SHA-512 fingerprint of the hateful content is computed on the device, together with the exact URL and time where it was found, and sealed with an RFC 3161/eIDAS timestamp anchored to the public blockchain inside an .asice container. The result is an immutable, admissible chain of custody. This proof does not delete the post or prevent the creation of hate, nor does it detect it; but it offers verifiable evidence the authorities can use to apply the law without depending on a manipulable screenshot.

The law is in place; the unalterable proof is missing

Since the European Union published the DSA guidelines on the protection of minors (July 14, 2025) [10], member states must push concrete measures against cyberbullying and grooming. In Spain, the draft law approved on June 4, 2024 [11] introduces the so-called "online restraining order" and reinforces the existing criminal framework: Article 510 of the Criminal Code punishes hate with 1 to 4 years' prison and fines, and Article 172 ter criminalizes repeated harassment with penalties of between three months and two years [12]. In the United Kingdom, the Online Safety Act of 2023, in force since July 25, 2025, obliges platforms to protect minors from bullying and hateful content, with fines of up to 10% of worldwide turnover [13]. Australia has gone a step further with the Online Safety Amendment (Social Media Minimum Age) Act of 2024, which bans social media access for anyone under 16 from December 10, 2025, with penalties of roughly 50 million Australian dollars for infringing companies [14].

All these regulations converge on one fundamental requirement: the need to **prove** what happened, when it happened and what takedown request was made. In practice, this requirement hits the same wall that has hindered the defense of victims for years. A screenshot, however detailed, does not guarantee the exact date; metadata can be manipulated or erased. A notice sent to the platform often vanishes without leaving any verifiable time trace, and the offensive content can be deleted before the authorities can gather solid evidence.

This fragility of the proof is not just a technical issue, but a legal obstacle that prevents the effective application of the laws: **the lack of reliable evidence sets aggressors free**, perpetuates the victims' silence and keeps the cycle of online hate and harassment alive.

As we have seen, **t2p** computes the SHA-512 fingerprint of the hateful content directly on the victim's device, without the material ever leaving the phone. To this fingerprint is attached the information of **where** (URL or identifier) and **when** (exact timestamp), together with the formal takedown request; the timestamp, based on the RFC 3161/eIDAS protocol and anchored to the public blockchain, creates a dated, verifiable chain of custody inside an .asice container [15]. This process produces a **neutral** proof, independent of any creator or intermediary, that can be presented to the courts as admissible evidence.

The proof of existence applied

In a Barcelona classroom, a secondary school teacher receives an email with a link to the capture of a chat where a group of students has sent hateful messages and false accusations against a 14-year-old classmate. The content is no longer available: the sender's phone shows only the notification "message deleted." Without tangible evidence, the school is trapped between the need to act and the impossibility of proving that the repeated harassment (Art. 172 ter) really took place. This is where **t2p** comes in: it computes on the device a SHA-512 fingerprint of the message, adds the link and the detection timestamp, and seals it all with a timestamp — as we have seen — anchored to the public blockchain inside an .asice container. The result is an unalterable chain of custody that can be presented as admissible evidence both in the school file and before the judge.

1. The school file

According to the report of the ANAR Foundation and Mutua Madrileña, that high percentage of cyberbullies who are classmates from the same school has already been noted [4]. These dynamics often translate into a "nothing gets done," leaving victims without recourse. With **t2p**, the student can produce digital proof that breaks this silence: the fingerprint shows that the message existed, when and where it was posted, and that its removal was requested. The school, obliged by the DSA guidelines to protect minors [10], can act immediately, open a formal investigation and impose disciplinary measures against the aggressor, without needing to keep the full content.

2. The criminal complaint

Article 510 of the Spanish Criminal Code punishes public incitement to hatred with up to four years' prison and fines [12]. To show that a series of posts constitutes "insistent and repeated" harassment (Art. 172 ter), temporal continuity and harmful intent must be proven. The chain of custody generated by **t2p** offers exactly that: each fingerprint includes the time and the place, and creates a chronological record that can be presented as direct proof of the repetitive character of the hate. Thus the victim does not have to store or re-expose the offensive messages; she submits only the fingerprint, which is immutable and admissible.

3. Notice-and-action for the platforms

The European DSA and the United Kingdom's Online Safety Act oblige social networks to remove illegal content within a reasonable time [10][13]. A sealed receipt generated at the moment of the report works as proof that the platform was notified and when. This prevents memory-holing: even if the aggressor deletes the post, the fingerprint and the date prove unequivocally that the content existed and that the required response was met.

Complementary cases

In cases like Pikesville's, where a synthetic audio defamed a school principal [8], provenance — the C2PA mark and the AI Act's labeling requirement [16] — is the complement that helps establish that a piece of content is synthetic, and t2p adds the neutral, dated proof of the harm. Moreover, the on-device personhood proof, as we have seen in chapter 108, makes it possible to verify "one person, one voice" without resorting to biometric deanonymization, and preserves freedom of expression and privacy.

In short, **t2p** does not eliminate hate or prevent new harmful messages from being generated; nor does it create a centralized registry of the abusive content. Its strength lies in providing an immutable fingerprint on the victim's device, which proves the existence and date of the harm without exposing the content itself [17]. This is the key piece that turns legal obligations into workable actions, and defends the digital dignity of the young against online hate.

The promise and the peril

In a world where any comment, meme or video can be captured, replicated and spread in seconds, the victim of digital hate often finds herself without solid proof of the harm she has suffered. The promise of **t2p** is to give this person power back: to generate, on her own device, a SHA-512 fingerprint of the offensive content, add the URL and the exact time, and register at the same time the takedown request. This timestamp, conforming to the specifications we have already cited and anchored to the public blockchain inside an .asice container, creates an immutable, dated chain of custody that can be presented to the courts as admissible evidence, with no need to store or re-expose the content. Thus the victim holds verifiable evidence of the harm and of the legal notice, capable of enforcing the obligations established against cyberbullying and Articles 510 and 172 ter of the Spanish Criminal Code [12].

Yet this technology carries a series of dangers that must be acknowledged with the same clarity as the promise. First, **it creates no centralized registry** of the abusive content: all the fingerprints are generated and kept solely on the victim's device, just as the StopNCII.org model does [17]. This prevents a global database of abuse from becoming an instrument of mass surveillance.

Second, **it can never serve to deanonymize or surveil** users. The seal proves the existence of a harm suffered by a person, but does not reveal exactly who they are; turning it into a tool to identify any voice would run against freedom of expression and threaten satire, whistleblowing and other legitimate forms of dissent (see chapter 62). Deanonymizing speech could freeze public debate and protect only the powerful; that is why the proof's value is, precisely, its neutrality.

A third danger is the **need for expiry or revocability** of the fingerprint. The immutability that guarantees judicial validity can also become a perpetual record of victimization, violating the right to be forgotten and perpetuating the psychological harm. That is why any implementation must allow the fingerprints to be erased or invalidated when the victim asks.

Finally, remember that **the seal is evidentiary infrastructure, not a cure**. The root of the problem lies in the attention-based business models that reward outrage, as shown by the "as we have seen" of hate speech on X after the change of ownership [9]. Without structural changes in these platforms — and without an honest debate on age verification, which protects minors but can also surveil them [10][14] — no technical tool will be more than a patch. The promise of **t2p** is powerful; the danger is trusting that technology can substitute for the legal and social reforms needed to guarantee dignity and safety to the young online. The seal is a tool of dignity and justice, not a technological shield that erases the harm.

It's time to act

Digital violence is no fleeting shadow: it leaves real scars on young people who live their adolescence between screens and notifications. 46% of American teenagers have suffered some form of cyberbullying [2] and half have endured it in the past year [5]; in Europe, one in six young people (16%) has already experienced some, with a rising trend since 2018 [3]. Laws like Law 71/2017, born of Carolina Picchio's case, have established clear penalties against cyberbullying and online hate, but the problem persists because the proof of the harm remains weak: manipulable screenshots, notices without a timestamp and content the aggressor deletes before it is documented.

For victims and families: the first step is to seal on the device a SHA-512 fingerprint of the offensive content, together with the URL and the exact time of detection. This fingerprint is combined with a formal takedown request and registered in an .asice container, with the fingerprint anchored to the public blockchain and a complementary RFC 3161/eIDAS timestamp [15]. Thus the victim keeps unalterable proof of the harm without having to store the offensive content, and avoids reliving it each time.

For schools: schools must treat the sealed fingerprint as a document of the file. Each registered case can be reviewed by multidisciplinary teams (psychologists, social workers and lawyers) who, on a verifiable chain of custody, can apply fast, legally sound measures: psychological support, sanctions or mediation.

For the platforms: the European DSA and the United Kingdom's Online Safety Act impose clear obligations; every notice must generate a sealed receipt, with date and time, delivered to the reporter and archived as a fingerprint only on the device. This mechanism complies with the DSA guidelines on the protection of minors [10] and avoids creating centralized databases that violate privacy.

For police and courts: the dated fingerprint must be accepted as admissible evidence for Articles 510 and 172 ter of the Spanish Criminal Code, since it constitutes an unalterable chain of custody proving the presence and content of the hate speech at the exact moment. Judges can use this evidence to issue takedown orders and impose sanctions without depending on proof that can be deleted or altered.

For legislators: the sealed proof must be legally recognized as a legitimate form of evidence, with a provision that guarantees its preservation and, at the same time, the right to be forgotten, so that the harm is not perpetuated but responsibility remains traceable.

As the case of Carolina Picchio reminds us [1], Law 71/2017 was born to give a voice to kids like her; now we can offer every young person a sealed proof of the harm and an enforceable takedown — without ever centralizing the content or forcing her to relive it. On this front that touches dignity, **time2act**.

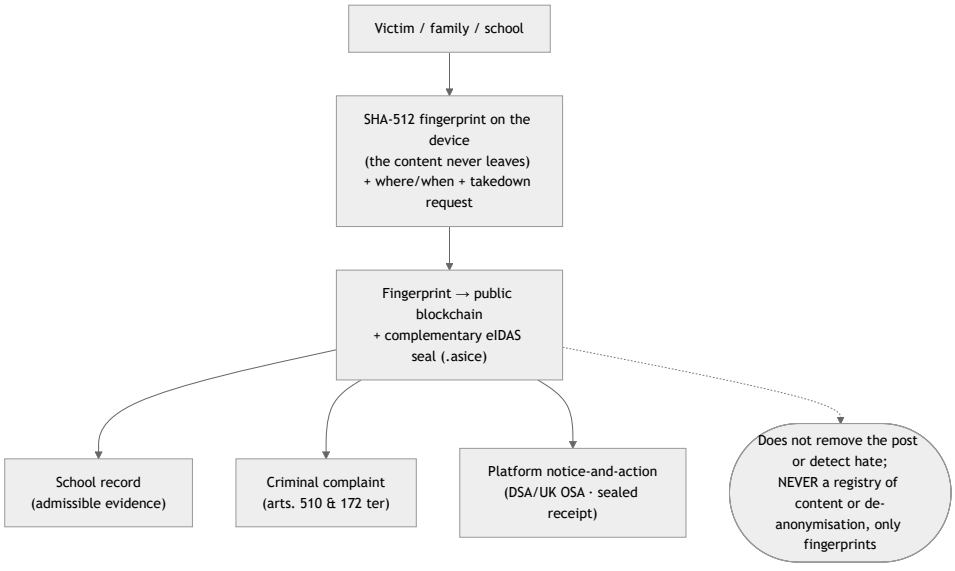


Figure 84.1 — The victim, the family or the school computes on the device the SHA-512 fingerprint of the hateful content (the content does not leave the phone) and adds where and when it was found and the takedown request; the whole's fingerprint is anchored to the public blockchain, with a complementary eIDAS seal (.asice). Three outcomes: the school file, the criminal complaint (Arts. 510 and 172 ter) and the platform's notice-and-action (DSA/UK OSA), with a sealed receipt. Limit: the seal does not remove the post, does not prevent the creation of hate and does not detect it; never a registry of the content or a deanonymization tool, only fingerprints on the device. Sources: the Picchio case; Molly Russell; ANAR; StopNCII.

---

[1] The Local (Italy) — *Italy unanimously passes anti-cyberbullying law* (May 2017; Law 71/2017; Carolina Picchio, 14, in 2013; Senator Elena Ferrara). https://www.thelocal.it/20170517/italy-unanimously-passes-anti-cyberbullying-law-online-abuse-

trolling ↩

[2] Pew Research Center — *Teens and Cyberbullying 2022* (December 15, 2022; 46%; name-calling 32%, rumors 22%, threats 10%). https://www.pewresearch.org/internet/2022/12/15/teens-and-cyberbullying-2022/ ↩

[3] WHO-Europe and the HBSC study — *One in six school-aged children experiences cyberbullying* (March 27, 2024; 16%; 44 countries; ≈280,000 young people). https://www.who.int/europe/news/item/27-03-2024-one-in-six-school-aged-children-experiences-cyberbullying--finds-new-who-europe-study ↩

[4] ANAR Foundation and Mutua Madrileña Foundation — *VIII Informe de Prevención del Acoso Escolar* (2025; 12.3% vs 9.4%; 14.2% with AI; 91.7% from the same school; 17,562 students). https://www.anar.org/crece-el-acoso-escolar-impulsado-por-el-ciberbullying-y-el-uso-de-la-ia-para-cometerlo/ ↩

[5] ADL — *Online Hate and Harassment: The American Experience 2024* (June 2024; teenagers 50%; 36% spilling into the physical world). https://www.adl.org/resources/report/online-hate-and-harassment-american-experience-2024 ↩

[6] CBS News / ITV News — *Molly Russell inquest* (September 30, 2022; coroner Andrew Walker; 2,100 of 16,300 posts; the world's first such ruling). https://www.cbsnews.com/news/molly-russell-suicide-death-coroner-social-media-instagram-pinterest/ ↩

[7] Candice Odgers — *The great rewiring: is social media really behind an epidemic of teenage mental illness?* (Nature, March 2024; a reply to Jonathan Haidt, *The Anxious Generation*, 2024). https://www.nature.com/articles/d41586-024-00902-2 ↩

[8] CNN — *Pikesville principal's voice on racist recording was an AI-generated deepfake* (Dazhon Darien / Eric Eiswert; January-April 2024). https://www.cnn.com/2024/04/26/us/pikesville-principal-maryland-deepfake-cec/index.html ↩

[9] PNAS — *X under Musk's leadership: substantial hate and no reduction in inauthentic activity* (2025; +≈50% hate; +70% "likes"; coordinated activity not reduced). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11819532/ ↩

[10] European Commission — *Guidelines on the protection of minors under the Digital Services Act* (July 14, 2025; cyberbullying, age verification; systemic risks = illegal hate speech). https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-protection-minors ↩

[11] La Moncloa — *Anteproyecto de Ley Orgánica para la protección de las personas menores de edad en los entornos digitales* (Council of Ministers, June 4, 2024; "online restraining order"; age verification). https://www.lamoncloa.gob.es/consejodeministros/resumenes/Paginas/2024/040624-rueda-de-prensa-ministros.aspx ↩

[12] Criminal Code (Spain) — Arts. 510 (hate crime; 1-4 years' prison) and 172 ter (harassment/stalking; aggravated by the victim's age). https://www.boe.es/buscar/act.php?id=BOE-A-1995-25444 ↩

[13] Ofcom / GOV.UK — *Online Safety Act 2023: protection of children codes* (duties in force July 25, 2025; bullying and hateful content; fines up to 10% of worldwide turnover). https://www.gov.uk/government/collections/online-safety-act ↩

[14] eSafety Commissioner (Australia) — *Social media age restrictions* (Act 2024; under-16s from December 10, 2025; the first country). https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions ↩

[15] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container). See chapters 13 and 15. ↩

[16] EU Artificial Intelligence Act — *Article 50: transparency obligations* (marking of AI-generated or manipulated content). https://artificialintelligenceact.eu/article/50/ ↩

[17] StopNCII.org — *How StopNCII.org works* (fingerprint generated on the device; the content never leaves; the fingerprints-only model). https://stopncii.org/how-it-works/ ↩

The Elderly: The Cloned Voice and the Invisible Abuse

"Are the old really human beings? Seeing how our society treats them, one may well doubt it."

— Simone de Beauvoir, The Coming of Age (1970)

An ordinary afternoon, in an ordinary apartment, an elderly mother's phone lights up: "Hi Mom, my phone broke, this is my new number. Message me on WhatsApp when you can." The text, simple and urgent, is the bait of one of the most widespread scams in Spain. The scammer has pulled the name and the family relationship from social networks or leaked databases and has fabricated a believable urgency — the recent variant says "my phone got wet"; INCIBE recommends always verifying through another channel, and the National Police has issued repeated alerts [1]. And this is only the written version. With today's artificial intelligence, the next call can carry the real son's cloned voice, fabricated from a few seconds of audio taken from a public video (see chapter 68) [3].

The figures from the American thermometer give the scale of what is coming. According to the FBI, in 2024 people over 60 lost close to $4.9 billion to cybercrime, 43% more than the year before; in 2025, complaints rose to 201,266 and losses to $7.75 billion (+59%), with more than 3,100 senior complaints mentioning AI — more than $352 million lost — and more than $5 million in distress scams with cloned voices [2].

This book devoted a chapter to digital harm against the young (chapter 84); this is its mirror for the elderly. Voice cloning technology is explained in chapter 68; wills and inheritance, in chapter 52; the life certificate, in chapter 64; corporate financial fraud, in chapter 16. The key question: how do we protect our elders — from the scam, from abuse and from pressure — without turning them into supervised minors? The answer is to give them, and their families, verifiable evidence with **t2p** — the SHA-512 fingerprint on the device, the RFC 3161/eIDAS timestamp and the anchor to the public blockchain [11] — under a principle Spanish law has already made its own and that admits no renunciation: support, not substitution (Law 8/2021 [10]).

The scam that calls home

Family impersonation has made the leap from text to voice. The FBI warned in December 2024 that generative AI is being used to clone relatives' voices in the so-called "grandparent" or distress scams: the call that cries and begs for urgent money no longer needs an actor — a few seconds of audio of the grandchild or the son, taken from a video posted on any network, are enough [3]. In 2025, declared losses from these distress scams with cloned voices exceeded $5 million, within a total of more than $352 million lost by seniors to frauds with an AI component [2]. Every call is an invasion of the deepest intimacy: the sound of a beloved voice, turned against the person who loves it most.

The romance scam completes the picture. The U.S. FTC received close to 60,000 complaints in 2024, with losses of $1.2 billion; and it is the 65-to-74 bracket that suffers the highest losses — more accumulated savings, less familiarity with digital tactics [4]. In Spain the scheme is fully active, with arrests announced by the Ministry of the Interior [6], and it has reached the darkest extreme: the Morata de Tajuña crime, in 2023 — three elderly siblings murdered after having been victims of a romance scam of up to €263,000 [5].

The common pattern of all these schemes is twofold. First, the scammer exploits the victim's best virtue: trust in family and in affection. Second, he counts on the best possible ally: the shame of reporting. When the victim stays silent, there is no proof, no statistic and no lesson learned. And there is still a third front, the most uncomfortable of all: when the danger comes not from outside but from within — from the one who cares, or the one who awaits the inheritance — proof becomes almost impossible. The next two sections go straight into it.

The abuse no one sees

In January 2022, the journalist Victor Castanet published in France *Les Fossoyeurs* ("The Gravediggers"), an investigation into Orpea, one of Europe's largest nursing home groups, with some 260,000 residents: rationing of food and hygiene products to cut costs in a group subsidized with public money, and medical

neglect. The French prosecutor's office opened criminal investigations into signs of mistreatment, misuse of public funds and labor violations, with dozens of complaints from families — some for involuntary manslaughter; the French government announced a sweeping inspection of the group's homes, and the state ended up injecting €605 million in exchange for a majority stake [7]. The lesson is structural: abuse in care happens behind closed doors, gets erased in the routine, and the word of a person with dementia does not weigh as proof. Without evidence, neither families nor inspectors can tell the exemplary home from the one that rations diapers.

Spain offers the counterpoint in small format, and with a way out. In a nursing home in El Escorial, the daughters of a resident with dementia — informed that their mother might be suffering mistreatment — installed a hidden camera. The worker was convicted, with a €480 fine and €3,000 in compensation upheld by the Provincial Court of Madrid, and the camera was endorsed as valid evidence [8]. What makes it possible is the Supreme Court's doctrine: with prior, well-founded suspicion of a serious crime, a temporary camera focused on the specific place passes the triple proportionality test — suitability, necessity and strict proportionality — and holds in criminal proceedings, even if administratively it may constitute a data protection infringement [9]. The nuance is essential: this is no open bar for spying on caregivers; it is a rigorous exception for well-founded suspicion.

And yet a gap remains: the homemade recording, without a certain date or attested integrity, is still the weak flank of any proceeding — "when was this recorded? who says it hasn't been edited?" It is exactly the gap the proof of existence closes, as we shall now see.

The last will under pressure

Chapter 52 dealt with the forged will: the fake signature, the altered document. Here the problem is subtler and more frequent: the authentic will of a pressured volition. The relative who isolates, the "caregiver" who suggests, the visit to the notary extracted on a bad day. No forgery anywhere — and yet no free will.

The Spanish legal framework made a paradigm shift with Law 8/2021, of June 2: from substitution to support [10]. The elderly or disabled person can make a will if at that moment they have the capacity to understand and express their volition; they can order voluntary measures before a notary — preventive powers, self-appointed guardianship, advance directives; whoever provides support must render accounts; and guardianship becomes exceptional. It is an architecture of dignity: the person decides, not the family.

But the dispute comes years later, often when the person is gone. Was there capacity? Was there pressure? The notary attests to the act — to that moment, before him — but the months of prior isolation or influence are seen by no one. The proof of existence supplies what was missing: temporal depth. One's own sealed conversations — the elderly person documenting, on their own device, the pressures they receive; the decision expressed and dated while still serene; the accounting of whoever provides support, anchored so no one can rewrite it [11]. The will stops being a snapshot of one instant and becomes a coherent line through time — much harder to capture, and much harder to dispute.

The proof of existence applied

The same technology that clones a son's voice can certify the real son's voice.

1. The real son

Against the cloned call, family verification. The family agrees on a safe word and a protocol: faced with any money emergency, hang up and call back the usual number [1]. And the real son can send a sealed proof of life: a short audio or video answering the agreed challenge, with the SHA-512 fingerprint computed on the device and the timestamp of the exact minute [11]. The cloned voice can imitate the timbre, but it cannot answer the challenge or fabricate a cryptographic seal of the present — it is the same challenge-and-freshness principle as the life certificate of chapter 64, brought down to the scale of the family (the phone network's authentication is dealt with in chapter 68). The house rule becomes simple: whoever calls asking for money and cannot seal the answer is lying — by default.

2. The family that watches with respect

Within the strict limits of the Supreme Court's doctrine — well-founded suspicion, a temporary, focused camera [9] — the recording a family is forced to make is worth far more if it is born sealed: each segment with its fingerprint and time anchored [11], an impeccable chain of custody from the first minute, immune to "this has been edited." In a case like El Escorial's [8], this armor would have spared any argument about the authenticity of the images. And the symmetry that makes it fair: nursing homes win too. Anchored care records — medication, hygiene, shifts, visits — protect the diligent professional from the unjust accusation and make detectable the systemic rationing that the Orpea case exposed [7] before it becomes a scandal. The proof has no side: it shelters the resident, the family and the good caregiver.

3. The attested will

An elderly person who wants to armor their last will — not against the family, but against doubt — can accompany the measures of Law 8/2021 [10] with anchored evidence: the video of the moment of decision, with a certain date; the conversations where they express this will over months, fingerprinted on their own device; the accounting of whoever provides support, anchored [11]. If someone ever alleges undue influence — or if the influence truly exists — there will be a verifiable chronology instead of interested memories. It does not replace the notary: it complements him with what no deed can contain — the consistency of a will over time.

4. The verified heart

Against the romance scam, the proof of existence cuts both ways. Whoever forms relationships online can ask the suitor for sealed proofs of life — whoever cannot attest even one verifiable minute of their real existence is suspect by default [11]. And the victim who begins to doubt can seal the conversations and the payment receipts on her device — without shame, without having to explain anything to anyone yet — so that, if one day she reports [6], she brings a complete chronology and not a handful of screenshots. The figures [4] and the memory of Morata de Tajuña [5] say clearly enough what is at stake.

It must also be said what the seal does not do: it does not detect the lie or the feigned love; it does not replace the Sunday phone call, or the company, or the decent funding of care — the Orpea case is a business model, not a cryptographic failure [7]. The proof of existence gives the elderly and their families what shame and the closed door took from them: the possibility of proving it.

The promise and the peril

The promise is a country where the fake son's call runs into a question it cannot answer [11], where abuse in care stops being word against silence [8], and where the last will reaches the courts attested over time and not reconstructed from memory [10]. But this chapter touches people in vulnerable situations, and the perils demand maximum frankness.

The first is **paternalistic surveillance**. Protecting cannot mean watching. If the family turns grandma's phone into a monitored device, we will have replaced the scammer with guardianship — exactly what Law 8/2021 wanted to banish [10]. The design rule: the tools are controlled by the elderly person — their device, their fingerprints, their decision to disclose. Support is provided with them, never over them.

The second is **the camera as the norm**. The Supreme Court's doctrine shelters an exception — well-founded suspicion, temporary, focused [9] — not the permanent video surveillance of caregivers, which would violate their rights and turn distrust into the standard of care. The seal makes lawful evidence solid; it does not make lawful a capture that was not.

The third is **the registry that must not exist**. Family verification works with challenges and fingerprints on each person's device; a central repository of "authentic" voices would be at once loot for the cloners and an infrastructure of mass surveillance. Only fingerprints, revocable, with no centralized biometrics and no deanonymization [11].

The fourth peril is no scam at all: it is putting technology where company is missing. No seal cures loneliness — which is the front door of the romance scam [4] and of the capture of the will. Real prevention is company, attention and well-funded care; cryptography only guarantees that, when proof is needed, it exists. If technology becomes an excuse to call home less, we will have lost far more than any scammer could take.

It's time to act

The generation that taught us to keep our word deserves tools to be able to prove it.

For families: agree on a safe word and the callback protocol [1]; faced with any money emergency, sealed verification before any transfer [11]; and talk about it openly at home — shame is the scammer's weapon.

For the elderly: your phone can be your certified memory. Seal the conversations that pressure you, the agreements that make you promises and your will when you have it clear [10][11]; disclosing it — or not — will always be your decision.

For nursing homes and caregivers: anchor the care records — medication, hygiene, shifts, visits: they protect the honest professional and make rationing visible before it becomes a scandal [7].

For notaries and jurists: combine the public attestation with temporal depth — the measures of Law 8/2021 accompanied by sealed chronologies of the will [10][11].

For banks and platforms: faced with unusual transfers by elderly customers, offer — never impose — the verification pause: a proof-of-life challenge to the alleged relative before executing the payment [3][11].

For legislators: recognize the capture-sealed proof in proceedings on abuse and on the capture of the will; and fund the inspection of nursing homes — no blockchain replaces an inspector [7].

The next "hi Mom" may carry your son's exact voice. Let your home's answer be a question only your son can answer — and a seal no machine can forge. For those who gave us our word, **time2act**.

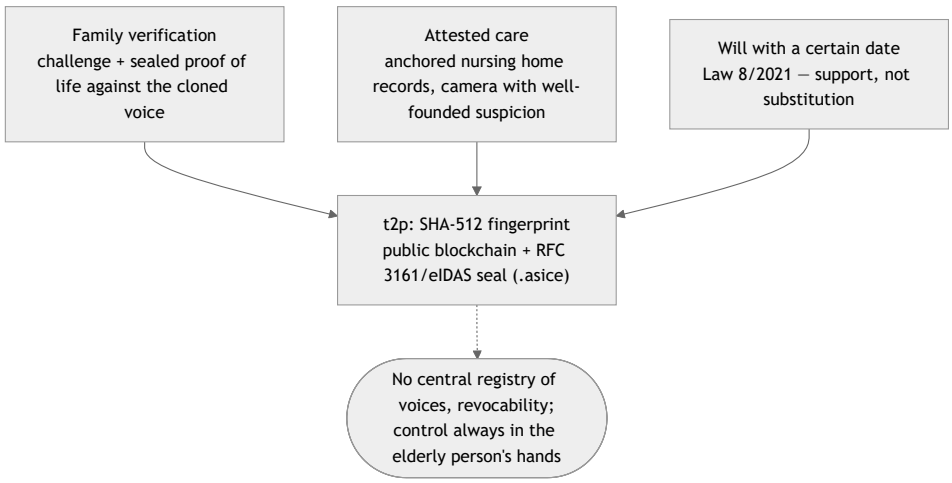


Figure 85.1 — Three shields for the elderly with a single logic (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, anchoring to the public blockchain, .asice): family verification against the cloned voice (challenge + sealed proof of life), attested care (anchored nursing home records; a camera only within the Supreme Court's proportionality doctrine) and the will with a certain date (Law 8/2021: support, not substitution). Limits: no central registry of voices, revocability, and control always in the hands of the elderly person. Sources: INCIBE; FBI/IC3; elDiario.es (El Escorial); Law 8/2021.

---

[1] INCIBE / Infobae — *El teléfono de tu hijo no se dañó: la estafa en WhatsApp que pide dinero para una emergencia* (May 8, 2025; the mechanics of the "hi Mom" scam, variants, verification advice); National Police (alert) — OKDiario; Valencia News — *Así funciona la estafa del «hijo en apuros» por SMS y WhatsApp*. https://www.infobae.com/tecno/2025/05/08/el-telefono-de-tu-hijo-no-se-dano-la-estafa-en-whatsapp-que-pide-dinero-para-solucionar-una-emergencia/ · https://okdiario.com/curiosidades/policia-nacional-alerta-sobre-esta-estafa-hola-mama-telefono-esta-roto-10582931 · https://valencianews.es/tendencias/mama-se-me-ha-roto-el-movil-asi-funciona-estafa-del-hijo-en-apuros-sms-y-whatsapp/ ↩

[2] HousingWire — *FBI: Seniors lost $7.75B to cybercrime in 2025 — a 59% jump* (201,266 complaints from over-60s; $4.9 billion in 2024, +43%; >3,100 senior complaints with AI and >$352 million; >$5 million in distress scams with cloned voices); PPC Land — *AI-linked fraud costs Americans $893 million* (22,364 AI complaints in 2025). https://www.housingwire.com/articles/fbi-seniors-cybercrime-2025/ · https://ppc.land/ai-linked-fraud-costs-americans-893-million-fbi-reports-for-2025/ ↩

[3] FBI / IC3 — *Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud* (PSA-241203, December 3, 2024; voice cloning in distress scams and family impersonation). https://www.ic3.gov/PSA/2024/PSA241203 ↩

[4] Factchequeado — *Qué son las estafas de romance y cómo identificarlas* (FTC: ≈60,000 complaints and $1.2 billion in 2024; the 65-74 bracket with the highest losses); AARP — *Los estafadores de romance engañan a sus víctimas con fraudes de criptomonedas*. https://factchequeado.com/teexplicamos/20240126/estafa-romance-scam/ · https://www.aarp.org/espanol/dinero/estafas-y-fraudes/info-2023/enganos-de-romance-con-criptomoneda.html ↩

[5] Wikipedia — *Estafa romántica* (the Morata de Tajuña crime, 2023: three elderly siblings murdered; a scam of up to €263,000). https://es.wikipedia.org/wiki/Estafa_rom%C3%A1ntica ↩

[6] Ministry of the Interior — *Detenidas dos personas por la estafa conocida como «Romance Scam», timo del amor*. https://www.interior.gob.es/opencms/es/detalle/articulo/Detenidas-dos-persona-por-la-estafa-conocida-como-Romance-Scam-timo-del-amor/ ↩

[7] Infobae — *Francia anuncia control de residencias de ancianos tras escándalo* (February 1, 2022; Victor Castanet's *Les Fossoyeurs*; Orpea ≈260,000 residents; rationing; criminal investigations into mistreatment, misappropriation and labor violations; involuntary manslaughter complaints); infoLibre — *El último escándalo de Orpea* (state injection of €605 million for a majority stake). https://www.infobae.com/america/agencias/2022/02/01/francia-anuncia-control-de-residencias-de-ancianos-tras-escandalo/ · https://www.infolibre.es/politica/ultimo-escandalo-orpea-negocios-trama-societaria-oculta-luxemburgo_1_1232899.html ↩

[8] elDiario.es — *Condenada la trabajadora de una residencia cazada con una cámara oculta mientras maltrataba a una interna* (El Escorial; a €480 fine and €3,000 compensation upheld by the Provincial Court of Madrid; the daughters' camera endorsed as evidence). https://www.eldiario.es/madrid/condenada-trabajadora-residencia-cazada-camara-oculta-maltrataba-interna_1_12498186.html ↩

[9] El Foro de Labos — *El Tribunal Supremo confirma la validez de una prueba obtenida con cámara oculta* (2022); Pardo Geijo Abogados — *¿Es legal poner cámaras ocultas en casa?* (well-founded suspicion of a serious crime; temporary, focused camera; triple proportionality test). https://www.elforodelabos.es/2022/09/el-tribunal-supremo-confirma-la-validez-de-una-prueba-obtenida-con-camara-oculta/ · https://pardogeijo.com/es-legal-grabar-camaras-ocultas-casa/ ↩

[10] BOE — *Law 8/2021, of June 2, reforming civil and procedural law to support persons with disabilities in the exercise of their legal capacity* (support, not substitution; voluntary measures before a notary; capacity to make a will; rendering of accounts; exceptional guardianship); El Notario — *Una visión práctica de la Ley 8/2021*. https://www.boe.es/buscar/act.php?id=BOE-A-2021-9233 · https://www.elnotario.es/practica-juridica/11767-una-vision-practica-de-la-ley-8-2021-de-2-de-junio-sobre-personas-con-discapacidad ↩

[11] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; proof of life with challenge; content never off the device; voluntary disclosure). See chapters 13, 15 and 64. ↩

Official Communications: Who Certifies That the Head of State Said This?

"Words had to change their ordinary meaning and to take that which was now given them."

— Thucydides, History of the Peloponnesian War, III, 82

In mid-June 2025, someone created a Signal account with the display name "marco.rubio@state.gov". With an AI-cloned voice and imitating the writing style of the U.S. secretary of state, **Marco Rubio**, the impostor contacted at least **three foreign ministers, a U.S. senator and a governor**: he left voice messages on Signal for two of those contacted and, in one case, sent a text message inviting the interlocutor to move to the same app. A State Department cable of **July 3, 2025** alerted employees: the likely goal was access to information or accounts, and the author has not been identified [3]. No system had to be hacked; a voice and a name in a messaging app were enough to pass off as good the word of American diplomacy.

The canonical precedent goes back to **March 16, 2022**, when attackers hacked the Ukrainian channel **Ukraine 24** and broadcast on it a deepfake of President **Volodymyr Zelensky** appearing to order soldiers to lay down their arms, while the hacked ticker claimed he had fled Kyiv; Zelensky refuted it within hours with a real selfie from central Kyiv ("a childish provocation") [1] — the case we analyzed in chapter 20. The question both episodes leave open runs through this chapter: when an official message can be forged and pushed through a trusted channel — or a personal one, like Signal — how can the public confirm quickly and independently what their government really said?

This chapter is distinct from chapter 48 — electoral integrity, where the target is the campaign, the candidate and the vote — and from chapter 20 — defense and human rights, where we spoke of battlefield command and proof of war crimes. Here the subject is the **authenticity of official crisis communications and of market-moving statements**: what a state or an institution says at the moment that matters most. The question is not whether a falsehood can be created — it can — but how an authority can publish its word so that any journalist or citizen can **verify it in seconds** as genuine, and how a forgery can be **refuted against the original** before panic does its work. It is this book's thesis applied to the most fragile moment of all: the emergency.

When the trusted channel becomes the weapon

The Ukraine 24 case is not anecdotal: as we have seen, it is the pattern. On **May 22, 2023**, an AI-generated image of an explosion near the **Pentagon**, spread by a fake account posing as "Bloomberg" and amplified by Russia's RT, briefly knocked the **S&P 500 down about 0.3%** (some 30 points) and the Dow about 80 points between 10:06 and 10:10 in the morning, recovering by 10:13; Bloomberg described it as "possibly the first time an AI-generated image has moved the market" [2]. The falsehood lived four minutes and moved billions: it did not need to be good, only to arrive through a channel that looked trustworthy.

In **April 2025**, Israel's **Channel 14** aired live, without verifying it, an Iranian-origin deepfake of former defense minister **Yoav Gallant** falsely claiming the U.S. could not defeat the Houthis; the presenter cut it mid-broadcast ("stop the video... this is fake") and the channel apologized [4].

From financial fraud to mass panic

The scammers' favorite branch is the figures who embody institutional trust. In February 2024, Mugur Isărescu, the world's longest-serving central bank governor, appeared in a deepfake promoting fraudulent investments — part of a wave that has hit central bank chiefs: Luxembourg's Gaston Reinesch, the governor of the Reserve Bank of India and, in June 2026, a forgery of the governor of the Bank of England, Andrew Bailey [5]. If the face that says "your money is safe" can be fabricated, the fraud does not just steal savings: it erodes the foundation of the whole system.

The other face is panic. As we have seen, on **January 13, 2018**, an official alert in Hawaii — not a deepfake, but a genuine human error — announced "BALLISTIC MISSILE THREAT INBOUND TO HAWAII. SEEK IMMEDIATE SHELTER. THIS IS NOT A DRILL," and the correction was delayed **38 minutes** [6]; a situation that could have generated mass panic and irreversible decisions had it been

forged. On **May 16, 2026**, the Singapore police revealed a Zoom video call with deepfakes of senior officials, including Prime Minister Lawrence Wong, that produced a minimum loss of **4.9 million Singapore dollars** ($3.8 million U.S.); Wong had already warned in March 2025 about the misuse of his image for scams [8].

The framework: authenticating at emission, but not at distribution

Regulators and security agencies have already acted, above all on the emission side. As we have seen, the European AI Regulation, Article 50 (Regulation 2024/1689), obliges deployers — including public authorities in matters of public interest — to disclose deepfake content (defined in Art. 3.60) from August 2, 2026, with penalties of up to €15 million or 3% of worldwide turnover [7]. In the U.S., FEMA's IPAWS alert system already digitally signs every official alert (EAS/WEA, CAP messages) and authenticates it via IPAWS-OPEN to prevent impersonation, with certificates renewed every three years: the baseline of authentication at emission. A cybersecurity information sheet from the NSA/CISA and their Five Eyes partners (January 29, 2025) recommends C2PA Content Credentials and durable provenance for government and critical infrastructure media, in line with eIDAS and qualified seals.

The hole, however, is obvious: this whole architecture protects the moment when the authority emits. It does not protect the moment when the message is distributed and rebroadcast — Ukraine 24's ticker, Channel 14's unverified rebroadcast, the fake "Bloomberg" image. It is on this stretch that the public is left defenseless.

Why the signature at emission alone does not save

It must be said bluntly. That IPAWS signs an alert or that a statement carries a qualified eIDAS seal solves authenticity at the origin, but it does not reach the citizen who receives the message forwarded on WhatsApp, sees it on a foreign TV news show or reads it on a hacked ticker. The signature lives in a closed system; the journalist who has to decide in two minutes whether or not to air a clip has no universal button to check it.

Worse, none of these controls exists at the place where the damage occurred. Ukraine 24 did not fail at Zelensky's emission — it failed because someone else pushed a falsehood through a channel the public believed its own. The signature at the origin only certifies what the authority emitted; it does not refute what an attacker has distributed. What is missing is a public check, vendor-neutral and offline, next to the distribution, that anyone can perform in seconds without asking anyone's permission.

The proof of existence applied

This is where **t2p** supplies the missing piece on the official artifact itself. The sealed object is the **authorized official statement at the moment of emission**: the press office's video or audio, the signed PDF statement, the central bank's decision or the payload of the emergency alert, before it leaves the issuing body. In the press office or the operations center, the capture application computes on the device the **SHA-512 fingerprint** — the file never leaves — obtains an **RFC 3161/eIDAS timestamp** and anchors the fingerprint to the public blockchain, producing an **.asice** container verifiable offline. For live statements by the principal — a head of state recording a denial, a minister at the lectern — a **proof of life** and the **forensic GPS** bind the act to that person, that place and that minute. The authority then publishes the **verification handle** (a QR or URL on the official channel) so that anyone can confirm in seconds that *this* exact file comes from the genuine source, at the declared time, bit-for-bit unaltered.

Thus a hacked ticker, an unverified rebroadcast or a forged clip can be **refuted against the sealed original** in minutes — like Zelensky's selfie, but cryptographically provable. It complements, not replaces, the controls at emission: the IPAWS signatures, the C2PA/Content Credentials provenance (as we have seen in chapters 13 and 15), the qualified eIDAS seals and the detection tools.

The seal proves **the origin, the time, the author and the integrity** of the official artifact — not that its content is true, legal or in the public interest. If an authorized official issues a false statement, in a panic or unconstitutionally — or makes an honest mistake like the 2018 Hawaii alert — a valid seal will prove that the wrong message really did come from the right office. It does not stop a deepfake that is never compared against the original, nor does it prevent the public from ignoring the verification handle;

nor does it replace editorial verification, institutional credibility or the existing layers of C2PA/SynthID, IPAWS signing and detection.

The promise and the peril

The promise is to close the distribution hole: to give the authority — and the public — a **fast, public and independent** way to separate the authentic official word from the forgery, at the moment of crisis when every minute counts. Zelensky's selfie worked because he is Zelensky and we could believe him; a sealed verification handle would work for any official, without depending on fame or improvisation, and with mathematical proof instead of blind trust.

The peril calls for nuances. The first: **false security** — "it is sealed, therefore it is true" — when the seal only guarantees the origin, not the truth or legality of the content. The second: if the issuing systems are hacked into signing a malicious payload, the seal would **authenticate the lie** with full solemnity; the issuer's security is a precondition, not a consequence. And the third: the **liar's dividend** worsens if the public learns that "unsealed" equals "false," because then any inconvenient unsealed statement can be dismissed as fake. The proof of existence is a layer of public verification, not a certificate of truthfulness or a substitute for judgment.

Moreover, the protocol's decentralized design can be a facade concealing a real concentration of power: who defines the update rules, and which nodes have the authority to sign them? If the community merely votes on proposals put forward by a small group of experts or corporate entities, the apparent consensus becomes a decision imposed from within. In this scenario, control of the code — and therefore of the sealing mechanism — passes into the hands of an elite that can adjust the conditions of emission to favor particular interests, such as suppressing inconvenient statements or introducing hidden filters. As Marc Goodman warns, "whoever controls the code will control the world, for better or for worse" [12]; the vulnerability lies not only in the technology, but in who holds the update key and how transparent their decisions are. The promise of public trust is therefore threatened if the protocol's governance remains opaque or centralized beneath a layer of apparent decentralization.

It's time to act

For **press offices and government cabinets**: seal every official statement, video and audio at the moment of emission and publish its verification handle on the official channel. For **emergency authorities** (FEMA and its counterparts): extend IPAWS-style signing with a public seal verifiable offline that the citizen can check without access to the closed system. For **newsrooms and news agencies**: demand and check the verification handle before rebroadcasting an official clip — the reflex Channel 14 lacked. For **central banks and financial institutions**: seal every market-moving statement, so that a deepfake of Isărescu or Bailey can be shown false in seconds. And for **regulators** (the AI Office, under Article 50): recognize the sealed handle as a means of transparency compliance and as admissible evidence of what the authority said and when.

In Kyiv, a president had to record a selfie to prove he was still at his post. Giving every authority — and every citizen — a way to prove cryptographically what was officially said, without depending on fame or luck, and never forgetting that the seal does not make what is declared in it true, is, on this front where the official word is the last line of defense, time2act.

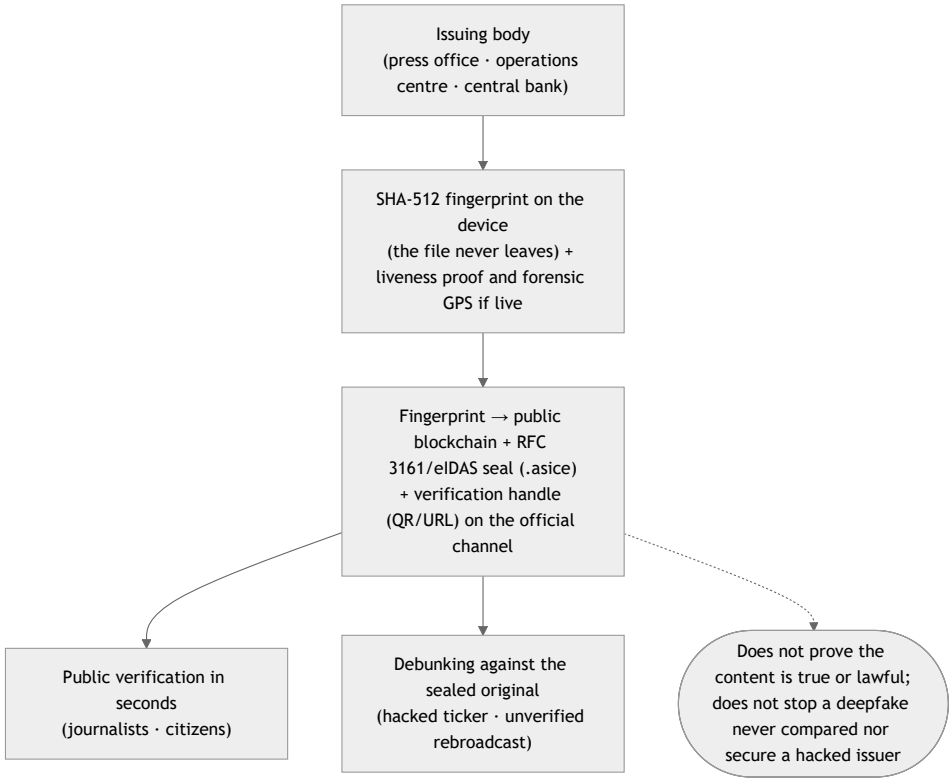


Figure 86.1 — The issuing body (press office, operations center, central bank) computes on the device the SHA-512 fingerprint of the official statement — video, release, decision or alert — without the file leaving; for live statements by the principal it adds proof of life and forensic GPS. The whole's fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), and its verification handle (QR/URL) is published on the official channel. Two outcomes: on the left, public verification in seconds by journalists and citizens; on the right, the refutation of a forgery (hacked ticker, unverified rebroadcast) against the sealed original. Limit: the seal proves origin, time, author and integrity, not that the content is true, legal or sensible; it does not stop a deepfake never compared against it, nor protect the issuer from being hacked. Sources: the Ukraine 24/Zelensky case; the fake Pentagon image; the Rubio case; the Channel 14/Gallant case; the 2018 Hawaii alert; Article 50 of the AI Regulation; IPAWS; NSA/CISA Content Credentials CSI.

---

[1] Atlantic Council — *Russian War Report: Hacked news program and deepfake video spread false Zelenskyy claims* (March 16, 2022; hacking of Ukraine 24's ticker and website; surrender deepfake; "childish provocation" selfie refutation). https://www.atlanticcouncil.org/blogs/new-atlanticist/russian-war-report-hacked-news-program-and-deepfake-video-spread-false-zelenskyy-claims/ ↩

[2] NPR — *Fake viral images of an explosion at the Pentagon were probably created by AI* (May 22, 2023; fake "Bloomberg" account, RT amplification, brief market drop). https://www.npr.org/2023/05/22/1177590231/fake-viral-images-of-an-explosion-at-the-pentagon-were-probably-created-by-ai ↩

[3] NPR — *Impostor uses AI to impersonate Rubio and contact foreign and U.S. officials* (July 2025; the secretary of state's cloned voice on Signal; State Department cable of July 3). https://www.npr.org/2025/07/09/nx-s1-5462195/impostor-ai-impersonate-rubio-foreign-officials ↩

[4] Clarity — *Iranian Deepfake of Israeli Defense Minister Airs Live* (April 2025; Channel 14's unverified rebroadcast of a Yoav Gallant deepfake; live cut and apology). https://www.getclarity.ai/ai-deepfake-blog/iranian-deepfake-of-israeli-defense-minister-airs-live----and-how-clarity-caught-it-instantly ↩

[5] Romania Insider — *Romanian central bank governor Mugur Isarescu targeted by deepfake* (February 2024; a fraudulent-investment deepfake; a pattern across several central banks). https://www.romania-insider.com/bnr-mugur-isarescu-deepfake-february-2024 ↩

[6] Wikipedia — *2018 Hawaii false missile alert* (January 13, 2018; false WEA ballistic missile alert; 38-minute correction; genuine error). https://en.wikipedia.org/wiki/2018_Hawaii_false_missile_alert ↩

[7] EU Artificial Intelligence Act — *Article 50: Transparency Obligations* (deepfake disclosure, including public authorities; applicable from August 2, 2026; definition in Art. 3.60; penalties of up to €15 million / 3%). https://artificialintelligenceact.eu/article/50/ ↩

[8] NSA/CISA et al. — *Content Credentials: Strengthening Multimedia Integrity in the Generative AI Era* (CSI, January 29, 2025; Five Eyes guidance endorsing C2PA Content Credentials for government and critical infrastructure media). https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF ↩

[9] FEMA — *Integrated Public Alert & Warning System (IPAWS)* (official alerts digitally signed by the issuing authority and authenticated by IPAWS-OPEN; the baseline of authentication at emission). https://www.fema.gov/emergency-managers/practitioners/integrated-public-alert-warning-system

[10] Singapore Police Force — *Footage from Zoom video conference involving impersonation of senior government officials* (May 16, 2026; deepfake of PM Lawrence Wong, President Tharman and MAS representatives; a 4.9 million Singapore dollar loss by a single victim). https://www.police.gov.sg/media-hub/news/2026/05/20260516_footage_from_zoom_video_conference_involving_impersonation_of_senior_government_officials

[11] Bulletin of the Atomic Scientists — *AI misinformation is threatening emergency communications. Here's how to fix that* (September 2025; false official alerts, the "liar's dividend" and the authentication of emergency communications). https://thebulletin.org/2025/09/ai-misinformation-is-threatening-emergency-communications-heres-how-to-fix-that/

[12] Marc Goodman — *Delitos del futuro* (Ariel, 2015). ↩

Impersonating the CFO: Who Really Authorized This Payment?

"It is prudent never to trust completely those who have deceived us even once."

— René Descartes, Meditations on First Philosophy, I (1641)

In January 2024, at the Hong Kong office of the British engineering and design firm **Arup**, an employee in the finance department received the instruction to transfer some **HK$200 million** — around **$25.6 million** U.S. At first he was suspicious: the email originating the order looked to him like a phishing scam. But he was invited to a **video conference** where the group's UK-based chief financial officer was present, along with several other colleagues; all of them had the exact face and voice of who they claimed to be. Reassured by what he saw and heard, the employee executed **15 transfers** to five Hong Kong bank accounts over the course of a week. Every participant on the call, except him, was an AI-generated deepfake. Arup's global chief information officer, Rob Greig, confirmed it months later: "we can confirm that fake voices and images were used," and that no internal system had been compromised — the fraud worked **purely by impersonating an authorization** [1]. As of 2025, no one had been arrested and the funds had not been recovered.

This chapter deals with a very specific face of fraud, distinct from chapter 16's — which covers the broad landscape of financial crime and regulatory compliance. Here we speak not of generic fraud detection, but of a single, discrete act: **the authorization of a payment or an executive instruction**. The question is not whether the face on the screen is real — we have already seen that is impossible to guarantee — but how to prove that **a directive really comes from the executive**. The book's thesis applies directly: instead of chasing the falsehood, we seal the truth at the origin.

A fraud that needs to break nothing

What is chilling about the Arup case is that no system was breached: there was no intrusion, no malware, no stolen passwords. The fraud exploited the one thing no firewall protects: **human trust in a familiar face and voice**. Greig was explicit that the internal systems were untouched; the weakness was procedural, not technical, because the company could not distinguish a genuine authorization from the chief financial officer from a fabricated one — since the authorization itself had never been sealed.

As we have seen, in March 2025 a chief financial officer of a multinational in **Singapore** authorized a payment of **$499,000** during a Zoom call in which the chief executive and the other executives were all AI-generated deepfakes; not one participant was real [7]. In **July 2024**, a **Ferrari** executive foiled a voice clone of the chief executive, Benedetto Vigna — pushed through an unknown WhatsApp number to sign, urgently, a confidentiality agreement on a secret acquisition — with a question only the real Vigna could answer: the title of a book he had recently recommended. The call cut off abruptly [5]. And in **May 2024**, scammers impersonated the chief executive of **WPP**, Mark Read, with a fake WhatsApp account, a Microsoft Teams meeting, a cloned voice and YouTube footage; the attempt failed with no losses, and Read warned staff that "we all need to be vigilant to the techniques that go beyond emails to take advantage of virtual meetings, AI and deepfakes" [6]. The difference between Arup and Ferrari was, literally, one question against the current.

The scale of the problem, in figures

Executive impersonation is no anecdote: it is a growing industry. **Business email compromise** (BEC) remains one of the costliest categories tracked by the FBI's complaint center (IC3): around **$3 billion** in declared losses, and almost **$8.5 billion** lost to BEC over the past three years [8]. And generative AI is pouring gasoline on it. Deloitte's **Center for Financial Services** projects that losses from generative-AI-enabled fraud in the United States could reach **$40 billion in 2027**, up from $12.3 billion in 2023 — a compound annual growth of **32%** — with email fraud and BEC potentially reaching **$11.5 billion** in the most aggressive scenario [2]. The Arup case, with its **$25.6 million**, is for now one of the largest documented deepfake-enabled payment frauds, but the trend indicates it will stop being exceptional.

What the regulators say

The authorities have taken note. On **November 13, 2024**, the U.S. Treasury's financial intelligence unit, **FinCEN**, issued alert **FIN-2024-Alert004**, warning that criminals use AI-generated deepfake media — fraudulent identity documents, audio and video — to defeat financial institutions' identity verification and authentication controls; it recommends **live checks** and **phishing-resistant multi-factor authentication** [3]. On **May 15, 2025**, the FBI's IC3 published advisory **PSA250515**, warning that malicious actors are sending AI-generated voice and text messages impersonating senior officials, and advises confirming identity **through previously verified channels** and using **pre-agreed verification phrases**, because AI-generated content "has advanced to the point that it is often difficult to identify" [4].

The European legal framework supplies the key piece: the **eIDAS Regulation (EU) No. 910/2014**, Articles 41 and 42, recognizes a **qualified electronic timestamp** with a presumption of accuracy and integrity that cannot be denied as evidence in any member state; it is the backbone for sealing an authorization artifact. In the United Kingdom, the **Solicitors Regulation Authority** updated its money laundering risk assessment on March 5, 2024, flagging deepfakes as an emerging risk for payment instructions and urgent client communications. The APP regimes in the UK and the EU already treat these fraudulent payments as a category of their own that demands stronger authorization controls.

Why the alert alone does not save

It must be said bluntly. FinCEN's and the FBI's recommendations are sensible — live verification, callback on a known channel, pre-agreed phrases — but they all depend on **someone, under pressure, remembering and applying a procedure** at the exact moment an apparently furious executive demands an urgent payment. The Arup employee had good instincts: he was suspicious of the email. What defeated him was the **social proof** of a room full of familiar faces. A callback policy is useless if the victim believes they are already talking to the boss. A pre-agreed verification phrase protects if you remember it and if the other party should know it too; the Ferrari executive was saved by instinct, not procedure. The underlying problem is that today authorization is a **subjective judgment** — "this looks like the boss" — instead of an **objective check**. As long as the question is "does this seem real to me?", deepfakes will always end up winning, because every year they seem more real.

The proof of existence applied

This is where **t2p** inverts the trust model. The artifact to protect is not the funds: it is the **authorization directive itself** — the executive instruction. Instead of the finance team trying (and failing) to detect whether the face on the call is real, the **genuine executive proactively seals every directive they authorize**. When a payment or transfer order is issued, the executive captures it on their device: a signed instruction text, the approval of a PDF invoice or a brief video or voice authorization with **proof of life** and **forensic GPS** confirming a living person in a known place. The device locally computes the **SHA-512 fingerprint** — the instruction never leaves the device — obtains an **RFC 3161 timestamp conforming to eIDAS**, anchors the fingerprint to the public blockchain and generates an **.asice** container verifiable offline.

From here on, the treasury moves not one euro if the instruction does not arrive with a valid, dated seal bound to the executive, verified against the public anchor. The question "did the boss really approve it?" stops being a subjective judgment and becomes a cryptographic check; a deepfake chief financial officer on a Zoom call simply cannot produce the real executive's prior cryptographic seal, so an "urgent" instruction without a seal is, by design, a red flag. This attacks Arup's failure mode, where a convincing video conference substituted for a verifiable authorization. The layer complements, never replaces, the existing controls: it coexists with the external-channel callback, dual signature or "four eyes," payment screening and the C2PA/SynthID provenance layer, as we have seen in chapters 13 and 15.

The seal proves that a specific directive was issued by a specific key holder at a given instant and has not been altered — it proves the act of authorizing, not that the authorization is legitimate or fraud-free. If the executive is coerced, bribed or deceived into sealing a fraudulent payment, or their device or signing key is stolen, the seal will certify a bad instruction; it does not verify that the beneficiary is legitimate, or that the invoice is real, nor does it prevent the laundering of funds. Nor does it stop an insider with a legitimate key or detect a deepfake on an unsealed call — it only makes the absence of the seal suspicious. Verifying the humans on the call itself — each participant's liveness handshake, before the microphones open — is dealt with in chapter 88. Recovering funds already transferred depends on a "no seal, no payment" policy, and

the layer reinforces callbacks, dual control and human judgment without replacing auditing, segregation of duties or a skeptical finance chief.

The promise and the peril

The promise is to make authorization **verifiable** instead of believable: to move from "it looks like the boss" to "the seal checks out." A treasury with a "no seal, no payment" rule makes the most perfect deepfake executive in the world useless, because it lacks the one thing it cannot fabricate — the real executive's private key. The Arup employee, tomorrow, would not have to decide whether a face is real; he would only have to check a seal.

The peril calls for the usual nuances. The first: **key security**. If the seal becomes the only gate, the executive's signing key becomes the target — it must be protected with hardware, biometrics and robust recovery, or the thief simply swaps the face for the key. The second: **false reassurance**. "It is sealed, therefore it is good" is a dangerous error; the seal certifies the origin, not the goodness of the instruction, and a coerced executive seals all the same. The third: **concentration and exclusion** — the rule must not shut out small suppliers or partners without the capacity to seal, nor become a barrier that favors only the big players. The proof of existence is a payment governance layer, not a certificate that the payment is sensible.

It's time to act

For executives and chief executives (Vigna, Read, finance directors): adopt the habit of sealing every payment authorization on the device at the moment of issuing it, so the treasury can demand its proof. For treasury and finance departments: impose the "no seal, no payment" rule and combine it with the callback and dual control; a valid seal must be a necessary condition, not a sufficient one. For banks and payment providers: offer seal verification as part of transfer screening, above all on large, urgent payments to new accounts. For regulators (FinCEN, FBI/IC3, SRA, the APP regimes): recognize the eIDAS-timestamped sealed directive as a valid authorization control and as evidence in fraud investigations. And for financial software and ERP vendors: integrate sealing and verification into approval workflows, so that "no seal, no payment" is the default path, not a heroic exception.

In Hong Kong, an employee with good instincts was defeated by a room full of faces that did not exist. Giving every authorization a sealed proof no deepfake can fabricate — never forgetting that the seal proves the act, not the goodness of the instruction — is, on this front where one click moves millions, time2act.

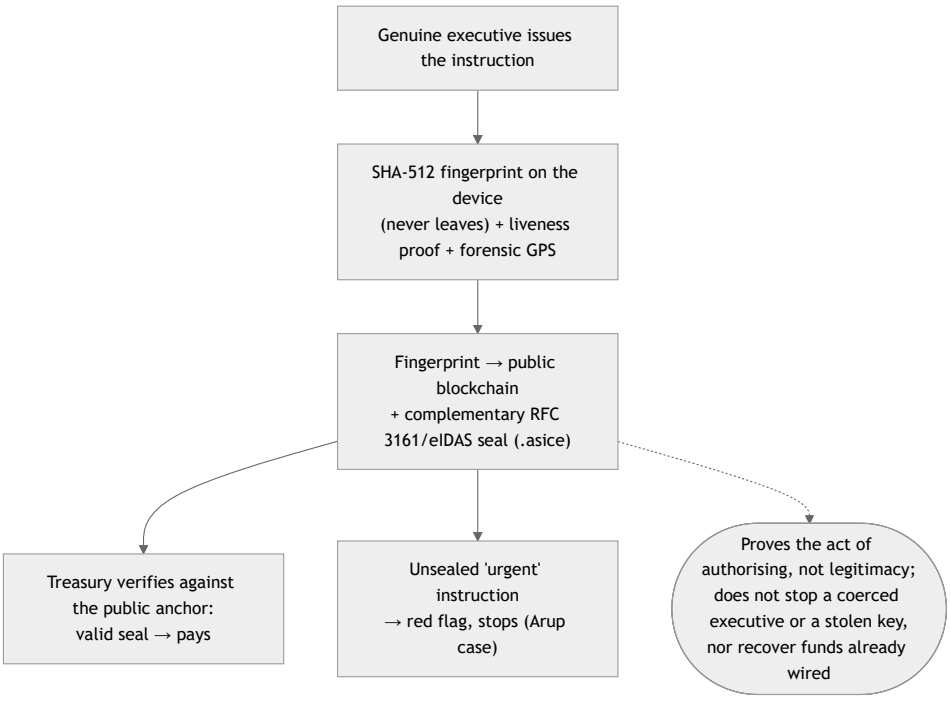


Figure 87.1 — The genuine executive captures the payment instruction on the device (signed text, PDF invoice or video with proof of life and forensic GPS), computes its SHA-512 fingerprint without it leaving the device, and anchors it to the public blockchain, with a complementary RFC 3161/eIDAS timestamp (.asice). The treasury verifies the seal against the public anchor: if it checks out, it pays; if it is missing, it stops — the "urgent" instruction without a seal is, by design, a red flag (Arup's failure mode). Limit: the seal proves the act of authorizing, not that the instruction is legitimate; it stops neither a coerced executive nor a stolen key, and does not recover funds already transferred. Sources: the Arup case; FinCEN alert FIN-2024-Alert004; FBI/IC3 PSA250515; eIDAS Arts. 41-42.

---

[1] Fortune — *A deepfake 'CFO' tricked British design firm Arup in $25 million fraud* (Arup confirmed as the victim; HK$200M/US$25.6M; 15 transfers; CIO Rob Greig's statement). https://fortune.com/europe/2024/05/17/arup-deepfake-fraud-scam-victim-hong-kong-25-million-cfo/ ↩

[2] Deloitte — *Generative AI is expected to magnify the risk of deepfakes and other fraud in banking* (projection: $40 billion in generative-AI fraud by 2027 from $12.3 billion in 2023; 32% annual growth; ~$11.5 billion in email/BEC fraud). https://www.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html ↩

[3] FinCEN — *FinCEN Issues Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions* (alert FIN-2024-Alert004, November 13, 2024; recommends live verification and phishing-resistant MFA). https://www.fincen.gov/news/news-releases/fincen-issues-alert-fraud-schemes-involving-deepfake-media-targeting-financial ↩

[4] FBI/IC3 — *Senior US Officials Impersonated in Malicious Messaging Campaign* (PSA250515, May 15, 2025; AI-generated voice/text; out-of-band verification and pre-agreed phrases). https://www.ic3.gov/PSA/2025/PSA250515 ↩

[5] Fortune — *Ferrari exec foils deepfake attempt by asking the scammer a question only CEO Benedetto Vigna could answer* (July 2024; a voice clone foiled with a control question). https://fortune.com/2024/07/27/ferrari-deepfake-attempt-scammer-security-question-ceo-benedetto-vigna-cybersecurity-ai/ ↩

[6] Marketing-Interactive — *WPP CEO Mark Read targeted by deepfake AI scam* (May 2024; fake WhatsApp, Teams, cloned voice and YouTube footage; failed attempt; Read's warning). https://www.marketing-interactive.com/wpp-ceo-mark-read-deepfake-ai-scam ↩

[7] Tookitaki — *The Deepfake Deception: How a CEO Video Call Cost $499K in Singapore* (March 2025; a finance director authorizes $499,000 on a Zoom call with every participant a deepfake). https://www.tookitaki.com/blog/deepfake-ceo-scam-singapore-2025 ↩

[8] Nacha — *FBI's IC3 Finds Almost $8.5 Billion Lost to Business Email Compromise in Last Three Years* (~$3 billion/year; ~$8.5 billion over three years; the payment fraud baseline). https://www.nacha.org/news/fbis-ic3-finds-almost-85-billion-lost-business-email-compromise-last-three-years ↩

[9] Trend Micro — *Deepfake CFO Video Calls Result in $25MM in Damages* (technical analysis of the multi-participant video call attack pattern used against Arup). https://www.trendmicro.com/en_us/research/24/b/deepfake-video-calls.html

[10] EUR-Lex — *Regulation (EU) No. 910/2014 (eIDAS), Articles 41-42* (legal effect of electronic timestamps: presumption of accuracy of date and time and admissibility as evidence; the legal basis of the sealed authorization artifact). https://eur-lex.europa.eu/eli/reg/2014/910/oj

[11] World Economic Forum — *Cybercrime: Lessons learned from a $25m deepfake attack* (February 2025; analysis of the Arup case and the control failures; the out-of-band verification argument). https://www.weforum.org/stories/2025/02/deepfake-ai-cybercrime-arup/

Live Identity: Who Is Really on the Other Side of the Screen?

"When an individual plays a part he implicitly requests his observers to take seriously the impression that is fostered before them."

— Erving Goffman, The Presentation of Self in Everyday Life (1959)

The setup WIRED documented in 2024 is chillingly simple: two phones on stands and a ring light. With the first, the scammer makes the video call — often on Zoom. This phone's rear camera does not point at him: it points at the screen of the second, which runs a face-swap app aimed at his face. What the victim sees is a face that blinks, smiles and answers in real time — and does not exist. The "relationships" stretch on for weeks or months, video calls included, until the request for money arrives. These are the "Yahoo Boys," the brotherhood of romance scammers operating mostly out of Nigeria — and they are brazen enough to show off the setups in Telegram groups, videos and all, sometimes to sell them to other scammers [1]. The social engineering expert Rachel Tobac summed it up: these apps have improved dramatically in a year — the skin tone, the hair, the eye color [1].

And it is no artisanal anecdote: it is an industry. iProov, a company that verifies remote identities for banking and governments, documented in 2024 a 2,665% increase in attacks with native virtual cameras — software that injects synthetic video as if it came from the device's camera, now the leading attack vector — a 300% rise in face-swap attacks aimed precisely at systems with liveness detection, and close to 24,000 users selling attack technology as a service; its threat catalogue counts 127 face-swap tools and 91 virtual cameras [2]. And the new vector is even cheaper: turning a single photo into animated video, in two steps [2].

For a century, seeing someone move and respond live was the definitive informal proof that they were real and were themselves. That proof has died, and with it the trust of two everyday moments: when someone enters — the sign-up on a platform — and when someone presents themselves — the video call, the interview, the first date. This chapter is about recovering them. The executive's instruction and their deepfake-filled call are dealt with in chapter 87; the identity of who you hire remotely, in chapter 91; banking KYC and financial crime, in chapter 16; the unique person of voting and participation, in chapter 108; the above-threshold approval, in chapter 80; the cloned voice on the phone, in chapter 68. This chapter's key question: how can a person prove they are real, present and themselves, at the exact moment it matters — and how can the other side check it without having to believe anyone? The answer: with **t2p**'s video-selfie signature — random liveness challenges, the device's biometric verification, SHA-512 fingerprint, RFC 3161/eIDAS timestamp, anchor to the public blockchain and a certificate anyone can verify [8].

Live is no longer proof

Three fronts, one failure: trust in the movement of a face. The romance front we have already seen [1]. On the labor front, the FBI warned in June 2022 — before the current explosion of generative AI — of a rise in complaints about the use of deepfakes and stolen personal data to land remote jobs: IT, programming, database and software positions, some with access to customer data, financial information and corporate systems [4]. The signals that then gave the deception away — lips out of sync with the sound, a cough that cannot be seen — were already subtle; every generation of tools erases them a little more [4][2]. And the executive front has its own chapter: the Hong Kong room where every face was fake and $25.6 million was authorized, in chapter 87.

Paper does not save you either. OnlyFake, the underground service exposed in February 2024, generated fake IDs and passports "with neural networks" for $15 — United States, Italy, China, Russia, Argentina, Czechia, Canada; the journalists of 404 Media passed the KYC check of the exchange OKX with one of these IDs, and OKX uses a professional third-party verifier [5]. The document alone is dead; the selfie alone, mortally wounded [2].

The industry itself admits it. Gartner predicted that by 2026, 30% of companies would no longer consider identity verification and authentication solutions reliable on their own, because of deepfakes against facial biometrics [3]. Injection attacks — synthetic video served directly into the camera stream — grew 200% in 2023, and the presentation attack detection standards, designed for the mask and the printed photo, do

not cover digital injection [3]. Translation: the sector that lives from verifying faces admits the board has been flipped on it.

The instinctive response is a race of detectors: guessing whether the face is fake. The book has already shown where that leads: the detector errs in both directions and the liar collects the dividend (chapter 15). The way out is different: stop asking "is this face fake?" and start demanding "prove you are a living human, present and willing to sign it — now." Not detecting the impostor: making it trivial for the real human to accredit themselves and expensive for the impostor. And letting anyone check the proof, without trusting any vendor.

Accrediting instead of detecting: what exactly the video selfie does

The gesture, step by step. The video-selfie signature is t2p's strongest signature. The person opens the document — a PDF: a sign-up form, a meeting sheet, a statement — and begins: the front camera records a short video while the app poses random liveness challenges — blink naturally, plus two more among turning the head, smiling or tilting it — to be completed in under 60 seconds; the device requires biometric verification — the fingerprint, or the PIN as an alternative — and, with everything passed, the person draws their handwritten signature on the screen. On certifying, the app generates two transactions on the public blockchain — one for the signature, one for the video — and a dedicated certificate: a one-page PDF with the document, the signature, the video evidence and the two transactions [8]. Only one person may be in the frame; if someone leaves or the time runs out, nothing remains.

What the proof contains. The certificate carries the file's SHA-512 fingerprint — computed on the device — the RFC 3161 timestamp conforming to eIDAS, the anchor to the public blockchain, a QR code and a forensic security check: GPS against spoofing, temporal coherence, device integrity and hardware attestation [8]. That last detail matters more than it seems: a virtual camera injecting synthetic video normally needs a tampered device, and the certificate puts that on record. The .asice evidence package — an ETSI standard — verifies fully offline, with open tools, without the app and without trusting the manufacturer [8]. And privacy by design: the video does not leave the phone — only the fingerprint travels to the chain [8]; sharing the evidence is always the signer's decision.

What it proves and what it does not. The sealed video selfie attests that a real, present person performed that act of signing at that instant — a living human answering unpredictable challenges, not canned video. What it does not do is state anyone's civil name: the "verified identity" the app shows is that of the account and the device's biometrics, not of a civil registry. Who the face in the video is, the receiver establishes — comparing it with the person they know, with the one on the screen, with the one in the account's founding selfie — or else the institutional layer that is coming: the European identity wallet, which every member state must offer by December 2026 [7]. The division of labor is clean: t2p supplies the evidence artifact; the receiver or the institution, the name. Civil identity as infrastructure — who you are before the state, with its promises and perils — is the subject of chapter 47.

Freshness. A month-old video selfie attests to nothing about today. That is why the signed document is the challenge: if it contains the context of now — the meeting of such a day, the sign-up with such a username, a phrase chosen a minute ago — the proof cannot be recycled (not-before), and the anchor sets the ceiling on the other side (not-after). It is the temporal sandwich the book already knows from commit-and-reveal (chapters 45 and 46), applied to the person.

The proof of existence applied

Three moments, three variants of the same gesture.

1. The meeting handshake

Before the video call that matters — the closing, the due diligence, the final interview, the board meeting — the host drafts the meeting sheet: a one-page PDF with the date, the time, the participants and a freshly chosen context phrase. Each participant signs it with their video selfie on their device and drops the certificate in the meeting chat [8]. One minute before the microphones open, each has from the others: a video with liveness challenges passed moments ago, a handwritten signature, a certain date later than the invitation — and a sharp face to compare with the little window on the call. Live face-swap can imitate a face inside Zoom [1][2]; what it cannot do with the same ease is produce the certificate of this signature: it would take passing random challenges with the swapped face before the native camera of a device whose

integrity is assessed in the certificate [8]. And if the doubt persists, the question against the current — the title of the book you recommended a week ago — remains free and lethal (chapter 87). The handshake accredits the humans in the room; the instruction that comes out of it, if it moves money, additionally demands its own seal — the "no seal, no payment" rule of chapter 87.

2. Signing up as a human

Registration on a platform — the marketplace, the social network, the dating app — is the other moment. The sign-up form as the challenge document: the person signs it with a video selfie and the platform receives a verifiable artifact — a real, present person signed up today, with this username — without keeping any biometrics: the video stays with the user, and on the chain there are only fingerprints [8]. It is the workable version of what academic research demands under the name of "personhood credentials": proving you are a real person — not an AI — without revealing any personal data; the proposal that researchers from MIT, OpenAI and Microsoft published in 2024, with in-person or institutional issuance and cryptographic verification [6]. Pseudonymous-but-human: the account can call itself whatever it wants; what it cannot be is a script that opens ten thousand. For the services that require the legal name — the bank, age checks — the European wallet arriving in December 2026 supplies the civil identity [7]. And the countermodel has a name too: scanning irises in exchange for cryptocurrency and concentrating the biometrics of half the planet in a single database — the book has spoken of it already (chapter 108). Here it is the reverse: the biometrics stay on each person's device.

3. The return: who comes back, and who calls again

Live identity is not just day one. The account that returns for a sensitive operation — changing the recovery email, emptying the wallet, selling what it holds — can re-accredit itself with a new video selfie over the operation's document: the receiver compares today's face with the account's founding selfie — continuity is a chain of dated videos, and the comparison is done by the verifier, never the app [8]. The "Mom, it's me, my number changed" that asks for money is answered with the gesture of chapter 85: a liveness challenge before moving one euro. And the first video date after weeks of chat — the Yahoo Boys' exact territory [1] — admits the same ritual as a board meeting: a one-minute handshake. Whoever loves you will do it smiling — literally: smiling is one of the challenges [8]; whoever does not exist, cannot.

The limit, stated clearly: the seal recognizes no faces and assigns no names — the comparison belongs to the receiver, human or expert; it is not a deepfake detector — it is a reversed burden of proof; it does not replace KYC where the law requires it — it complements it with an artifact the customer controls; and not everything calls for a handshake: the everyday call with someone you have known all your life is not a due diligence. The usual criterion: the higher the risk, the stronger the gesture — the threshold of chapter 80.

The promise and the peril

The promise: meetings where every seat comes accredited by a certificate any participant can verify without trusting anyone [8]; sign-ups where creating ten thousand fake accounts stops being free, because each one demands a living human answering challenges before a camera; scammers who, to try, must leave a dated, signed artifact bound to an account — evidence working for the victim from minute zero; and the arriving institutional infrastructure — the European wallet [7] — completed by a neutral evidence layer that no government and no platform holds in custody.

Liveness is attackable too — honesty before marketing: iProov documents face-swaps aimed precisely at systems with liveness detection [2], and t2p's manual promises no anti-deepfake shield [8]. Gesture challenges are not a wall: they are a bar — higher in a native app, with the device's integrity and hardware attestation in the certificate [8], than in any webcam — but no bar is infinite. The answer is the whole book's: layers — the certificate, the human face comparison, the question against the current (chapter 87), the alternative channel — and asymmetric cost: whoever clears the bar leaves a dated, signed proof of their own fraud. The deepfake that signs gets notarized.

Friction — two credits and a minute per person is not for every call or every sign-up. It is for the moments that are worth it: the high amount, the first contact, the irreversible operation. Each sets their own threshold — and chapter 80 has already shown how.

Biometrics as a toll — the book's ethical line does not move: the proof of life is an option for whoever wants to accredit themselves, never a universal toll. The video stays on the device; on the chain, only fingerprints [8]; there must be alternatives for those who cannot or will not — the PIN already is one

within the gesture [8] — and anonymity remains legitimate and vital where showing your face kills (chapter 113).

Exclusion — whoever has no phone with biometrics, whoever cannot pay credits, whoever is seventy-eight and trembles before a timed challenge: if live accreditation becomes a condition of access, we have built a new door for leaving people out. Every platform that adopts it needs an equivalent in-person or assisted path — and not a second-class one.

The registry that pursues — a world where every sign-up and every meeting leaves an anchor is also a correlatable world: who got accredited, when, to talk to whom. The anchors must not be linkable to one another, disclosure must be selective — the disputed stretch, never the whole life — and control of each field must rest with the signer (chapter 112).

It's time to act

Let "I am real and I am me" become provable again.

For meeting and video call platforms: give the handshake a native place — a waiting room where certificates are attached and checked before the microphones open; let "all participants accredited" be as visible as the browser's padlock. Detection can help; accreditation decides.

For consumer platforms — networks, marketplaces, dating apps: offer the sign-up as a human — verifying the video-selfie artifact without holding any biometrics [8][6] — and flag the accounts that have passed it; without making it an obligation or penalizing legitimate anonymity.

For companies: a handshake policy for video calls above a threshold — remote hiring, the new supplier, the big deal — dovetailed with chapter 87's rule: the handshake accredits the humans in the meeting; the instruction's seal, the authorization that comes out of it. Two layers, two different questions.

For citizens: before sending money to someone you have only seen through a screen — the video date that turns urgent, the "son" with a new number, the recruiter with the perfect job — ask for a minute of handshake with liveness challenges [8]. Whoever is real passes it smiling; whoever refuses or gets angry has already answered you.

For regulators and standards bodies: the European wallet must arrive on time [7] and coexist with the humanity evidence of neutral layers the citizen controls; and liveness standards must cover the injection attack, not just the mask and the photo [3] — the hole is there.

The screen taught us to see each other from afar; deepfakes have taken away our right to believe what we see on it. It can be recovered: not by guessing who is fake, but by giving every human the way to prove they are real — here, now, signing this. Let every "it's me" — at the sign-up, in the meeting room, on the night someone you love asks for help — be able to carry its proof. So that we can look at the screen and know who is there, **time2act**.

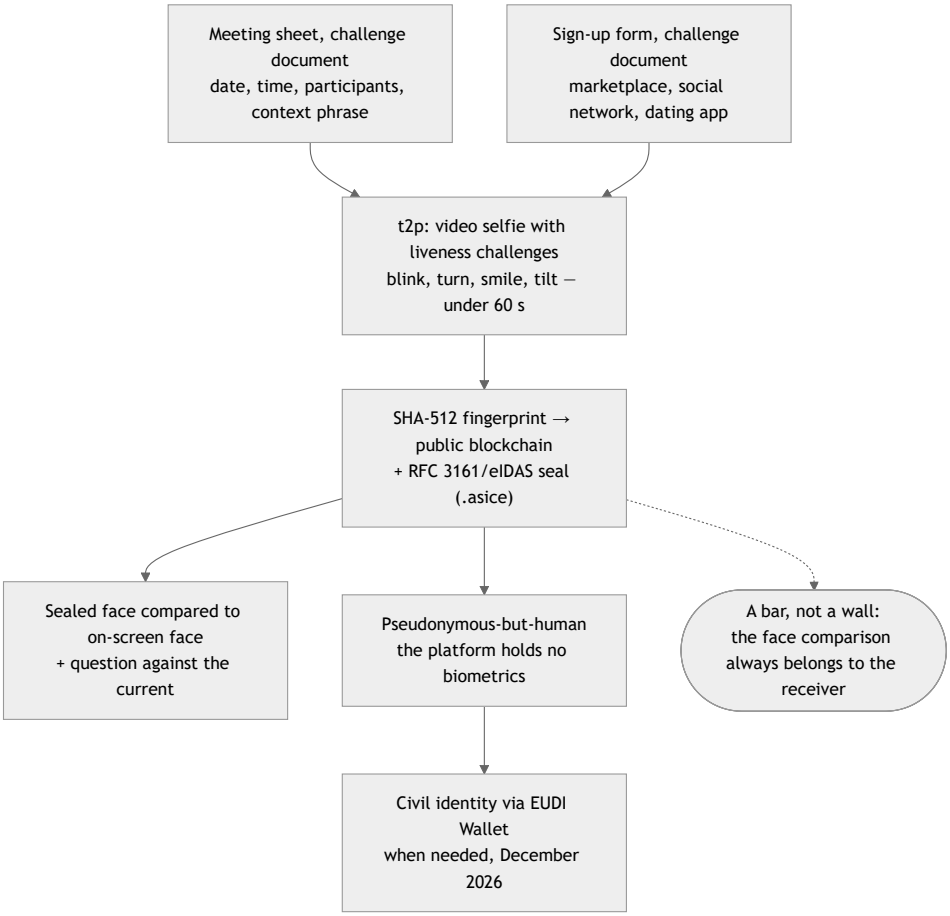


Figure 88.1 — Live identity with proof of existence (t2p: video selfie with random liveness challenges — blink, turn, smile, tilt — in <60 s, device biometrics or PIN, SHA-512 fingerprint, RFC 3161/eIDAS seal, anchor to the public blockchain, certificate with the video evidence and hardware attestation, .asice package verifiable offline): on the left, the meeting handshake (the meeting sheet as a challenge document signed by each participant before the microphones open; comparison of the sealed face with the on-screen one; the question against the current as an extra layer); on the right, signing up as a human (the sign-up form signed with a video selfie: pseudonymous-but-human, the platform verifies without holding biometrics; civil identity via the EUDI Wallet when needed). Context: real-time face-swap commercialized (Yahoo Boys), +2,665% native virtual camera attacks and +300% face swaps against liveness systems (iProov 2025), 30% projected enterprise distrust (Gartner). Limits: gesture challenges are a bar, not a wall — layers and asymmetric cost (the fraud that signs gets notarized); the face comparison belongs to the receiver, never the app. Sources: WIRED/9to5Mac; iProov; Gartner; FBI IC3; Decrypt; arXiv/MIT; EC (EUDI).

---

[1] 9to5Mac (on WIRED's reporting) — *Latest romance scam tactic uses two phones with face-swap app* (Apr-2024: the "Yahoo Boys" and real-time face-swap inside video calls — two phones on stands and a ring light, or a laptop with a webcam; they show off the setups on Telegram and sell them; Rachel Tobac: the apps have improved dramatically in a year). https://9to5mac.com/2024/04/18/romance-scam-face-swap/ ↩

[2] iProov — *Threat Intelligence Report 2025: Remote Identity Under Attack* (native virtual camera attacks +2,665%; face swaps +300% versus 2023, aimed at systems with liveness detection protocols; ~24,000 users selling attack technologies as a service; image-to-video conversion as the new vector; 127 face-swap tools, 91 virtual cameras and 10 emulators catalogued). https://www.iproov.com/reports/threat-intelligence-report-2025-remote-identity-attack · https://www.iproov.com/press/annual-identity-verification-threat-intelligence-report ↩

[3] Gartner (via FutureCIO) — *Deepfakes to render identity verification and authentication solutions inefficient* (Feb-2024 prediction: by 2026, 30% of companies will no longer consider identity verification and authentication solutions reliable in isolation, due to deepfakes against facial biometrics; injection attacks +200% in 2023; presentation attack detection standards do not cover digital injection; Akif Khan: combine PAD, IAD and image inspection). https://futurecio.tech/deepfakes-to-render-identity-verification-and-authentication-solutions-inefficient/ ↩

[4] FBI / IC3 — *Deepfakes and Stolen PII Utilized to Apply for Remote Work Positions* (PSA of 28-Jun-2022: deepfakes and stolen personal data used to apply for remote IT, programming, database and software jobs — some with access to customer data and corporate systems; signals: lips and sound out of sync, coughs or sneezes that do not match the image). https://www.ic3.gov/PSA/2022/psa220628 ↩

[5] Decrypt — *People Are Using Basic AI to Bypass KYC — But Should You?* (Feb-2024: OnlyFake, fake IDs and passports generated "with neural networks" for $15 — U.S., Italy, China, Russia, Argentina, Czechia, Canada; 404 Media passed the KYC of

the exchange OKX, which uses the third-party verifier Jumio; users used them to open bank accounts and unlock crypto accounts). https://decrypt.co/216188/ai-generated-fake-id-bypass-kyc-aml-banks-crypto-onlyfakes ↩

[6] Adler, Hitzig, Jain et al. (arXiv) / MIT News — *Personhood credentials: Artificial intelligence and the value of privacy-preserving tools to distinguish who is real online* (Aug-2024; 32 authors, with researchers from MIT, OpenAI and Microsoft: digital credentials to prove you are a real person — not an AI — without revealing any personal data; in-person or institutional issuance; cryptographic verification; "AI is moving very fast, certainly much faster than the speed at which governments adapt"). https://arxiv.org/abs/2408.07892 · https://news.mit.edu/2024/3-questions-proving-humanity-online-0816 ↩

[7] European Commission / Corbado — the European Digital Identity Regulation (EU) 2024/1183: member states must offer an EUDI Wallet to citizens, residents and businesses (deadline: December 2026); fewer than a third of states meet the readiness bar today (Politecnico di Milano observatory). https://ec.europa.eu/digital-building-blocks/sites/spaces/EUDIGITALIDENTITYWALLET/pages/915931811/The+European+Digital+Identity+Regulation · https://www.corbado.com/blog/eudi-wallet-2026-deadline-rollout-eic-2026 ↩

[8] t2p — product manual (video-selfie signature: "the app's strongest signature" — the front camera records while you sign, a natural blink and two random liveness challenges (turn the head, smile, tilt it) in under 60 seconds, mandatory device biometric verification with PIN as an alternative, handwritten signature, 2 credits and two transactions; dedicated certificate with the document, the signature, the video evidence and the two transactions; only one person in the frame; certificate with SHA-512, RFC 3161/eIDAS seal, anchor to the public blockchain, QR and forensic check with device integrity and hardware attestation; ETSI .asice package verifiable offline without trusting the manufacturer; files stay on the device — only the fingerprint travels; the proof of life confirms "a real, present person" at signing). See chapters 13 and 15; executive impersonation and the "no seal, no payment" rule, in chapter 87; who you hire remotely, in chapter 91; commit-and-reveal, in chapters 45 and 46; the unique person of participation, in chapter 108. ↩

Information That Moves Markets: Who Said It, and When They Said It

"Rumor: no other evil is swifter; it thrives on movement and gains strength as it goes."

— Virgil, Aeneid, IV, 174-175

On November 10, 2022, at 1:36 in the afternoon, an impostor bearing the **$8** Twitter Blue "verified" badge (@EliLillyandCo, with Lilly's logo) posted: "We are excited to announce insulin is free now." The real stock of **Eli Lilly** fell from about $368 to about $346 in twenty minutes, momentarily erasing some **$15 billion** in value and forcing the company to apologize and point investors to its authentic account, @LillyPad [1]. Six months later, the same mechanics found a bigger target: on May 22, 2023, the AI-generated image of an explosion next to the **Pentagon** — the case that opened chapter 17 — amplified by "verified" accounts posing as "Bloomberg News" (@BloombergFeed) and by the Russian state account RT, knocked the S&P 500 down about 0.26-0.3% to its session low — the Dow lost some eighty points between 10:06 and 10:10 — before recovering fully within minutes, once the Arlington County fire department and the Pentagon confirmed there had been no explosion [2]. It was, according to Bloomberg itself, possibly the first time an AI image had moved the market.

This chapter is not about annual accounts or auditing — that is chapter 55's terrain — nor about the advertising fraud of chapter 56. Here the artifact is not a retrospective balance sheet, but a **single live disclosure**, a news item or a rumor that is published and moves a security's price in a matter of seconds. The question is simple and brutal: in a market where algorithms react to a headline in milliseconds, **how can the market know, instantly, whether a "news item" really comes from the issuer that claims to be its source?** And the book's thesis fits directly: against forgery, the proof of existence does not judge the truth, but it fixes **who** said it and **when**.

A lie faster than any denial

The Lilly case and the Pentagon case share a root with an older, more brutal precedent. On April 23, 2013 — the "Hack Crash" — a hacked Associated Press account posted that there had been explosions at the White House and that President Obama had been injured. The Dow fell 143.5 points and more than **$136 billion** was wiped off the S&P 500's value within seconds, before correcting [3]. The lesson is that **algorithmic and high-frequency traders** read the text of social networks and react to it in milliseconds: long before any human — let alone any company — can deny it. The asymmetry is total; publishing a false disclosure takes seconds, while correcting it depends on an authority or on the company itself verifying and denying it through a channel the market recognizes.

It is not always a prank that dissolves in minutes. In 2025, in the Ostin Technology scheme (NASDAQ: OST), deepfake videos of Goldman Sachs strategist David Kostin, of Elon Musk and of Mark Zuckerberg, combined with stolen identities of FINRA-registered advisers, inflated the stock **1,175%** — from $0.78 to $9.40 between mid-April and late June — before a **94% collapse in a single day**. The U.S. Department of Justice filed charges for securities fraud, wire fraud and conspiracy in September 2025; the scheme allegedly evaporated more than **$950 million** [4]. Here the deepfake is no longer a fleeting headline, but the engine of a complete pump-and-dump.

The rumor as a weapon, and the fraud of the legitimate channels

The problem is not only the impersonation of listed issuers. It is also the impersonation of those who give **advice**. The Bombay Stock Exchange (**BSE**) has had to issue repeated warnings — on January 12, 2026, again on March 8, and a fourth incident toward April — over deepfake videos of its chief executive, **Sundararaman Ramamurthy**, giving supposed stock tips and promising "extraordinary profits." The exchange has had to clarify that neither he nor any BSE executive makes recommendations or runs WhatsApp or Telegram groups, and has promised legal action [5].

And there is still a third vector, the most unsettling for a publisher: the **legitimate channels** can be poisoned too. A fake press release distributed through **GlobeNewswire** claimed Walmart would accept the cryptocurrency **Litecoin**; Litecoin's price rose more than **30%** before Walmart denied it [7]. If even

the regulated wires — the ones the market treats as authorized — can carry a forged disclosure, the question of authentic origin stops being theoretical.

What the law says about the disseminator

The law already pursues this conduct, but from the other end. In Europe, the Market Abuse Regulation (MAR) (Regulation EU No. 596/2014, Article 12) expressly includes within "market manipulation" the dissemination of information through the media — including the internet — that gives or is likely to give false or misleading signals about the supply, demand or price of a financial instrument, including rumors, where the disseminator knew or should have known it was false. Article 17 of the same MAR obliges issuers to disclose inside information to the public **as soon as possible** and through official channels; the Listing Act amendments, applicable from June 5, 2026, adjust its delay regime.

In the United States, Rule 10b-5 and Section 9(a)(2) of the Securities Exchange Act, enforced by the SEC, cover false statements and price-moving schemes, with parallel DOJ charges for securities fraud and conspiracy — as in the Ostin case. Reg FD and Form 8-K define the legitimate channels of corporate disclosure: EDGAR plus the recognized wires (Dow Jones, Bloomberg, Business Wire, PR Newswire). In a joint alert of January 25, 2024, the SEC, FINRA and NASAA warn that fraudsters may use deepfake videos to imitate a company's chief executive announcing false news with the intent of manipulating a stock's price [6]. As we have seen, the enforcement arsenal is forceful: in fiscal year 2025, the SEC opened 456 actions and ordered $17.9 billion in penalties for market manipulation and disclosure fraud [8].

Why the law alone does not save

It must be said bluntly: MAR and Rule 10b-5 punish **afterward**, when the harm is done and the price has already moved. They pursue whoever spreads the lie, but they give the market no way of knowing, **in the second the headline appears**, whether it really comes from Lilly or from an $8 impostor. The legitimate channel — EDGAR, the 8-K, the regulated wire — is exactly what the impersonation exploits: no one can check, instantly and online, that this tweet or this release has the origin it claims. The failure is one of **verifiable authentic origin**, and no amount of after-the-fact penalties closes it. As long as the market cannot tell the genuine from the fake in milliseconds, the algorithm will react first and ask questions later.

The proof of existence applied

This is where **t2p** supplies the missing piece, not as an oracle of truth, but as a **cryptographic reference of the genuine article**. When an issuer's investor relations team — or its filing agent, or the wire — publishes a price-sensitive item (a Form 8-K, an inside-information announcement under MAR Article 17, a chief executive's statement, a quarterly result, an official social post or a video), the artifact is sealed **on the device**: the **SHA-512 fingerprint** is computed locally — the file never leaves — wrapped in an **RFC 3161 timestamp conforming to eIDAS**, anchored to the public blockchain and issued as an **.asice** container verifiable offline. When an executive appears on camera making the statement, **forensic GPS** and a **proof of life** are added. The issuer publishes the verification fingerprint next to its official account and its investor relations page, so that any journalist, exchange, regulator, broker or algorithmic news channel can confirm in milliseconds whether a "breaking" tweet, screenshot or release really comes from the company — or is the impostor: the fake @EliLillyandCo tweet, the fake Bloomberg account of the Pentagon image, the BSE executive's deepfake video, Ostin's stolen identities. It replaces nothing: it sits **beneath** EDGAR/8-K, Reg FD, MAR and the regulated wires as an **authenticity layer**, and it complements the C2PA/SynthID content provenance of chapters 13 and 15 — the seal proves "this comes from the issuer," the provenance marks "this is AI-generated," and the regulator still judges the abuse.

Now the limit, in all its rawness. "Garbage in, garbage on-chain": the seal proves that a specific file was issued by an authenticated source at an instant with its integrity intact; **it does not prove that the sealed disclosure is true**, complete, legal or non-manipulative. A company can perfectly seal a release that is itself a lie — a pump-and-dump release sealed by the fraudsters' own shell company is authentic-but-false — and the seal neither stops an issuer that wants to deceive nor catches the disclosure that misleads by omission: that remains the job of the SEC, ESMA, FINRA, exchange surveillance and the auditors. Nor does it make an unsealed rumor "fail": it only lets verifiers confirm that the genuine one does not match it — and that depends on issuers actually sealing their real disclosures and on the market

checking them. It does not by itself prove that a video is real and not AI (that is C2PA/SynthID and detection), it does not assess materiality or insider intent, and it does not prevent the algorithmic reaction in the milliseconds before verification can happen. It is a layer of authenticity and certain dates, not an oracle of truth.

The promise and the peril

The promise is to offer the market a **genuine reference standard** against which to check any viral copy: had Lilly been able to publish the fingerprint of its real release, any verifier would have seen in a second that the impostor's tweet did not match. The honest issuer regains control of its own origin; the journalist, the broker and the news channel gain an instant check before amplifying.

The peril calls for the usual nuances. The first, **false security**: "it is sealed, therefore it is true" is precisely the confusion between integrity and truth that a sealed pump-and-dump would exploit. The second, **uneven coverage**: if only the big issuers seal, the absence of a seal at a small, honest company could be read, unjustly, as suspicion. And the third, **speed**: until verification is natively integrated into the algorithmic news channels, the seal will arrive late for the first millisecond wave. The proof of existence is an authenticity layer, not a substitute for human or regulatory judgment.

Moreover, the fact that the proof-of-existence seal becomes a source of massive data for predictive models creates a risk of **power over collective behavior**. Whenever an entity can anticipate which contents will be deemed authentic or false, it acquires the capacity to influence — or even shape — users' decisions before they act. As Carissa Véliz warns, "every attempt to predict human behavior is also an attempt to control it, which, at its extreme, leads to authoritarianism" [11]. If the algorithms processing the fingerprints concentrate in the hands of a few corporations or governments, the information about which messages are "verified" can be used to restrict freedom of expression and impose official narratives. That is why the registry's design must guarantee a decentralized, transparent architecture, with governance mechanisms that prevent the accumulation of predictive power in any single entity, ensuring that the authenticity layer keeps serving collective trust and not authoritarian surveillance.

It's time to act

For **issuers and their investor relations teams** (like Lilly after its own case): seal every price-sensitive disclosure at the moment of issuing it and publish its fingerprint next to the official account. For **exchanges and regulators** (SEC, ESMA, FINRA, BSE, SEBI): recognize the sealed origin as a verification signal and integrate its check into market surveillance. For the **regulated wires** (Business Wire, PR Newswire, GlobeNewswire): seal every release at origin to close the gap that allowed the Walmart-Litecoin fraud. For the **media and the algorithmic news channels** (Bloomberg, Dow Jones, aggregators): check the fingerprint before amplifying a "breaking" headline. And for **brokers and platforms** (X and the rest of the networks where the rumor is born): expose the verification status next to financial content, so the $8 blue badge is never again mistaken for an authorized source.

As we saw on May 22, an image that showed nothing real moved the S&P 500 before anyone could deny it. Offering the market a way to check, in the same second, who said it and when they said it — never pretending the seal says whether it is true — is, on this front where milliseconds are worth millions, **time2act**.

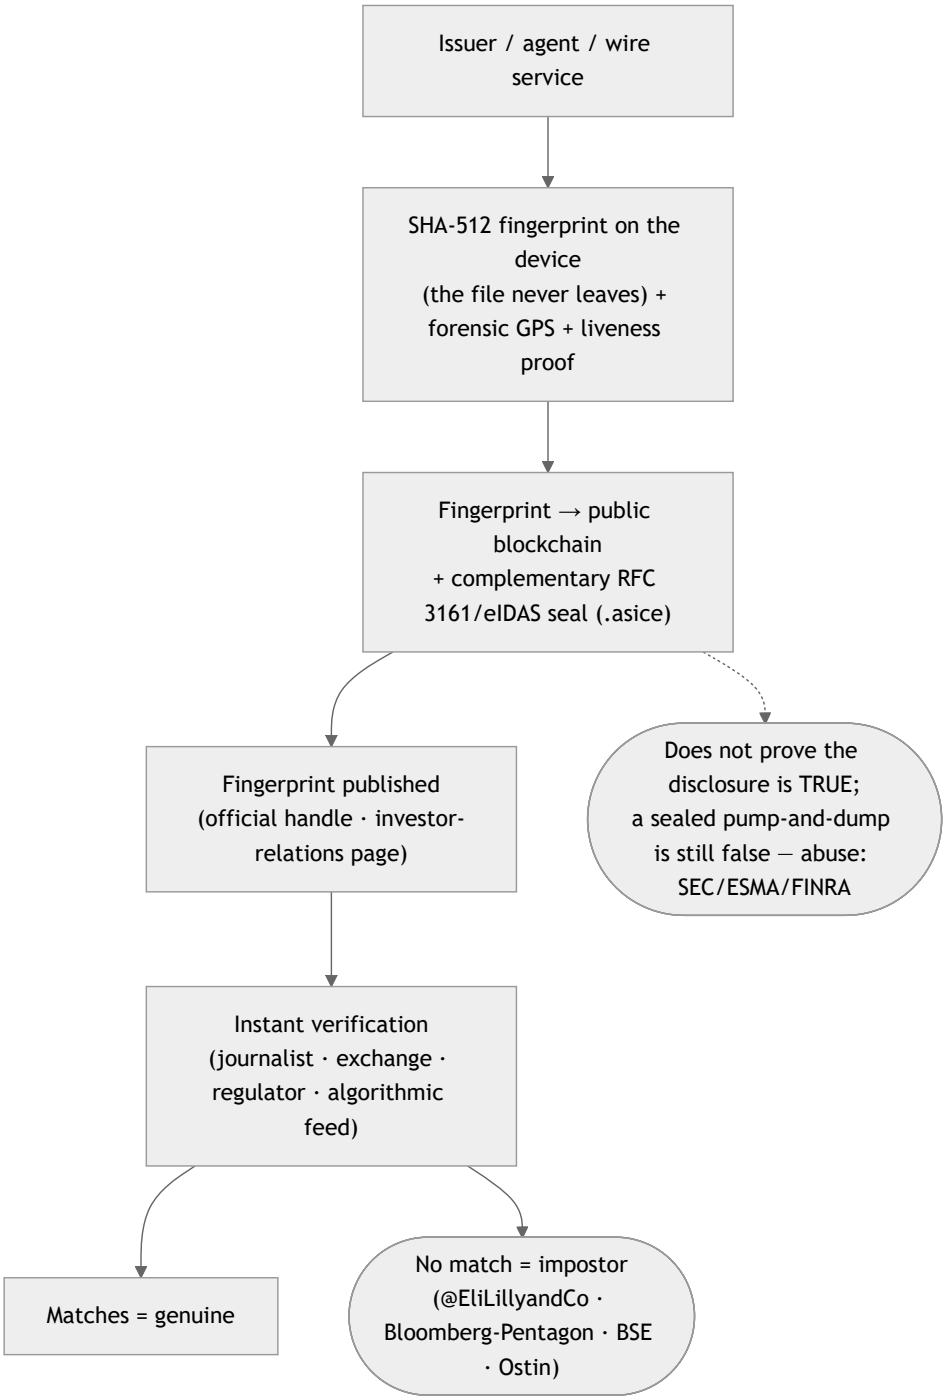


Figure 89.1 — The issuer (or its agent or wire) seals the price-sensitive disclosure at the moment of issuing it: SHA-512 fingerprint on the device → public blockchain, with a complementary RFC 3161/eIDAS timestamp (.asice), with forensic GPS and proof of life if an executive is on camera. The fingerprint is published next to the official account; any journalist, exchange, regulator or algorithmic channel confirms in milliseconds whether a "breaking" tweet or release matches the genuine one (Lilly, Pentagon-Bloomberg, BSE, Ostin) or is the impostor. Limit: the seal proves who issued it and when, not that the disclosure is true; a sealed pump-and-dump release is still false, and manipulation remains the job of the SEC/ESMA/FINRA and the exchanges. Sources: the Eli Lilly case; the fake Pentagon image; the 2013 Hack Crash; the Ostin Technology case; the SEC/FINRA/NASAA alert; MAR Art. 12.

---

[1] 340B Report — *Fake Lilly Twitter account tweets 'Insulin Is Free Now,' temporarily wiping billions from company's stock value* (Nov. 10, 2022, 1:36 p.m. ET; fall from ~$368 to ~$346; ~$15 billion erased; the company's apology and the @LillyPad account). https://340breport.com/fake-lilly-twitter-account-tweets-insulin-is-free-now-temporarily-wiping-billions-from-companys-stock-value/ ↩

[2] NPR — *Fake viral images of an explosion at the Pentagon were probably created by AI* (May 22, 2023; spread by verified accounts including a fake Bloomberg; brief S&P 500 drop; denied by the Arlington fire department). https://www.npr.org/2023/05/22/1177590231/fake-viral-images-of-an-explosion-at-the-pentagon-were-probably-created-by-ai ↩

[3] University at Buffalo — *With one false tweet, computer-based 'Hack Crash' led to real panic* (April 23, 2013; hacked AP tweet; Dow −143.5 points; >$136 billion erased in seconds; algorithmic and high-frequency reaction). https://www.buffalo.edu/news/releases/2015/05/037.html ↩

[4] InvestmentNews — *Shareholders accuse Ostin Technology of deepfake-driven $950M pump-and-dump scheme* (2025; deepfakes of Kostin/Musk/Zuckerberg; stolen identities of FINRA advisers; 1,175% inflation; 94% one-day collapse; DOJ indictment Sept. 2025). https://www.investmentnews.com/regulation-legal-compliance/legal-shareholders-accuse-ostin-technology-of-deepfake-driven-950m-pump-and-dump-scheme/265308 ↩

[5] The Hans India — *BSE warns investors about fake deepfake video misusing CEO's identity* (2026 warnings over deepfake videos of chief executive Sundararaman Ramamurthy; recurring incidents January-April; promised legal action). https://www.thehansindia.com/business/bse-warns-investors-about-fake-deepfake-video-misusing-ceos-identity-1038779 ↩

[6] SEC / FINRA / NASAA — *Investor Alert: Artificial Intelligence (AI) and Investment Fraud* (Investor.gov, Jan. 25, 2024; explicit warning about deepfake videos of chief executives to "manipulate a stock's price"). https://www.investor.gov/introduction-investing/general-resources/news-alerts/alerts-bulletins/investor-alerts/artificial-intelligence-fraud ↩

[7] PRNEWS — *Hoax Releases Hit Newswire Services* (fake Walmart-Litecoin release via GlobeNewswire; Litecoin +30%; the wire-authentication angle). https://www.prnewsonline.com/hoax-releases-hit-newswire-services/ ↩

[8] Sidley Austin — *SEC Enforcement FY2025 Results Signal Shift in Priorities* (456 actions; $17.9 billion ordered; adjusted $1.4 billion + $1.3 billion; priority on market manipulation and disclosure fraud). https://www.sidley.com/en/insights/newsupdates/2026/04/sec-enforcement-fy2025-results-signal-shift-in-priorities-in-direct-critique-of-prior-administration ↩

[9] EUR-Lex / Lexparency — *MAR Regulation (EU) No 596/2014, Article 12 — Market manipulation* (dissemination of false or misleading information and rumors through the media and the internet as market manipulation). https://lexparency.org/eu/32014R0596/ART_12/

[10] FINRA — *Artificial Intelligence (AI) and Investment Fraud* (FINRA's version of the alert; language on deepfake videos of chief executives and fake news to move a stock's price). https://www.finra.org/investors/insights/artificial-intelligence-and-investment-fraud

[11] Carissa Véliz — *Prophecy: Prediction, Power, and the Fight for the Future* (2025). ↩

Beneficial Ownership: Who Is the Flesh-and-Blood Person Behind the Company

"Respect is tied to names. Anonymity and respect are mutually exclusive."

— Byung-Chul Han, In the Swarm (2017)

On October 3, 2021, the **International Consortium of Investigative Journalists** (ICIJ) published the **Pandora Papers**: a leak of **11.9 million documents** — 2.94 terabytes — from **fourteen** offshore service firms. Inside were the true owners of more than **30,000 companies** and more than **29,000 beneficial owners** hidden behind chains of shell companies, trusts and foundations spread across more than two hundred countries. The files named **35 sitting or former heads of state and government** and more than **330 politicians**, among them King Abdullah II of Jordan, the then Czech prime minister Andrej Babiš and the presidents of Kenya, Ecuador and Gabon, plus more than **130 billionaires** — 46 of them Russian oligarchs [1]. It came five years after the **Panama Papers** of April 2016 — 11.5 million documents from a single firm, Mossack Fonseca, which shut down in 2018 [2] — and identified roughly twice as many beneficial owners. The common thread: the owner on paper, a nominee director or a layer of anonymous companies, was almost never the ultimate beneficial owner, and no registry of the time could prove **who** that person was, **at what moment** they had taken control, or whether the ownership declaration was **authentic**.

This chapter is not about the public contract of chapter 43 — tenders and awards — nor the invoice of chapter 54 — documentary fraud in trade — but about what lies **upstream** of any transaction: proving the **ultimate beneficial owner** behind the legal person, whatever that company later buys, sells or finances. The question is surgical: how do we fix **who** declares to be the owner, **what** they declare and **when**, so that the declaration can be neither backdated nor silently manipulated? It is the book's thesis applied to the heart of corporate opacity: the proof of existence does not uncover the secret, but it prevents denying afterward what was declared before.

An architecture of secrecy that survives every leak

The Pandora Papers were not an anomaly, but a confirmation. None of those companies was illegal in itself; what made them dangerous was that no one could prove **who** was hiding behind them. Ten years after the Panama Papers, the balance is sobering: tax investigations had recovered some **$1.36 billion** and triggered deep reforms — the American Corporate Transparency Act, the EU's beneficial ownership registers, the United Kingdom's Register of Overseas Entities — and even so Transparency International concludes bluntly that "**the fundamental architecture of offshore secrecy remains intact**" [3]. The reason is structural: every reform has demanded **declaring** who the owner is, but almost none has been able to **verify** it or **date** it unalterably.

The fraud, then, is not only the omission: it is the **false declaration**. A kleptocrat puts a front man at the helm, interposes a trust in an opaque jurisdiction, and the national registry obediently records a name that is not his. When a leak finally erupts — Panama, Paradise, Pandora — the damage has been running for years, and the judicial question is always the same: was this declaration like this on the day it was filed, or has it been touched up since? Without an independent temporal anchor, the answer depends on a registry that the declarant himself, or an accommodating official, may have modified.

Registries that declare themselves and are not checked

The technical problem has a name: **self-declaration without verification**. Open Ownership, the reference authority on beneficial ownership data, is blunt: most registries rest on data the obliged party itself declares, and **without verification** — compliance checks, cross-checks and review of the supporting evidence — the registries allow false information and front-man structures that hide real control [11]. A database field, mutable and unsealed, is not proof; it is an assertion.

The rollout figures illustrate it. The American **Corporate Transparency Act** initially estimated some **32 million** companies obliged to report, plus some 5 million new ones each year; by November 2024 only 6.5 million filings had been made [5]. And in March 2025 an interim FinCEN rule turned the whole regime upside down: it exempted **all** U.S. domestic companies and U.S. persons, and cut the covered population

from some 32 million to **fewer than 12,000** entities — the foreign ones [4]. The most ambitious registry in the world, before consolidating, was left practically empty. In Europe, meanwhile, the **Court of Justice of the EU** judgment of November 22, 2022 struck down general public access to beneficial ownership registers as a "**serious interference**" with the rights to private life and data protection (Articles 7 and 8 of the Charter): access is now limited to authorities, obliged entities and whoever shows a legitimate interest [6]. The result is a mosaic where the data exists, but its **reliability** and its **accessibility** remain in dispute.

A regulatory tide demanding accuracy

The 2024-2027 regulatory response is the most ambitious in history on this front. In Europe, **Regulation (EU) 2024/1624** — the single anti-money-laundering rulebook (AMLR) — will apply directly to the 27 states from **July 10, 2027**, and the new **Anti-Money Laundering Authority (AMLA)**, seated in Frankfurt, has been operational since **July 1, 2025**; its mandate includes **interoperable** beneficial ownership registers at EU scale [7]. In the United Kingdom, the **Economic Crime and Corporate Transparency Act 2023** has been rolling out since **November 18, 2025** the **mandatory identity verification** for directors and persons with significant control; between March 2024 and March 2025 alone, Companies House removed false or misleading information affecting **100,400 companies** and rejected more than **10,200 suspicious filings** [8]. In the United States, on **December 16, 2025** the **Eleventh Circuit**, in *National Small Business United v. Treasury*, became the highest American court to uphold the constitutionality of the CTA, though the March rule limits its practical scope [9]. And above it all, the **FATF's Recommendation 24**, revised in March 2022, requires every country to ensure that competent authorities have **adequate, accurate and up-to-date** beneficial ownership information through a multi-pronged approach, including a registry or equivalent mechanism [10]. This whole architecture asks for something paper documents do not give on their own: **verifiable, dated accuracy**.

Why the registry alone does not save

It must be said without detours: a beneficial ownership registry, however mandatory, **is only as reliable as the will of whoever fills it in and the capacity of whoever checks it**. A database entry can be edited; a declaration filed today can be backdated tomorrow without a trace; and an official can accept a name he knows to be false. The FinCEN case proves it in reverse: an entire regime can be emptied by one administrative rule, and the data that remains stops being comparable with yesterday's. A PDF of an ownership declaration is a PDF. What is missing is not the form or the law that requires it — both already exist — but an **unalterable anchor fixing who declared what, and at what exact instant**, so that neither the declarant nor the registry's custodian can alter it afterward without it showing.

The proof of existence applied

This is where **t2p** supplies the missing layer, and it does so on the declaration artifact, not on the corporate structure it will never see. When a beneficial owner, a director or a corporate service provider files or updates a beneficial ownership declaration, the seal binds **that specific declaration and its evidence** to a fixed point in time and identity. On filing, the declaration document plus the underlying proof — the shareholder register extract, the trust deed, the disclosure of the nominee agreement, the passport or ID page, and a **proof of life** of the declarant captured in the field — is fingerprinted with **SHA-512 on the declarant's device**; the fingerprint receives an **RFC 3161 timestamp conforming to eIDAS** and is anchored to the public blockchain, producing an **.asice** container verifiable offline. When a real person asserts control — the flesh-and-blood human behind the shell — the field capture adds **forensic GPS** and proof of life, showing that a living human, not a deepfake or a recycled photo, made the attestation in that place and at that moment.

This complements the registries and verification regimes — FinCEN's BOI, the AMLR/AMLA interoperable registers, Companies House identity verification, the FATF's R.24 — without replacing them: the registry still holds and discloses the data, and the gatekeeper still screens it, but now every declaration carries a tamper-proof, independently checkable receipt of who declared what, when, and that the declaration and its evidence have not been altered since. This turns a self-declared, mutable database entry into a sealed, court-proof attestation, and makes the silent backdating or later manipulation of an ownership record detectable. It is the complement, on the act of declaring, of what chapters 13 and 15

propose on content provenance with C2PA: there, what is produced is sealed; here, who declares to be the owner.

Now the limit, in all its rawness, because deceiving oneself here would do harm. "Garbage in, garbage on-chain": the seal proves the act of declaring, who did it, when, and that the declaration and the attached evidence are unaltered; **it does not prove that the declared person is the ultimate beneficial owner**. A kleptocrat can stand before the camera, pass the proof of life and seal a perjured attestation naming a front man, or omit the hidden trust entirely, and the seal will faithfully date the lie. It cannot pierce a chain of foreign companies it never sees, nor does it detect a nominee agreement, nor does it know whether a notarized shareholder register is already false at the origin. It does not verify identity documents against the issuing authorities, does not assess whether the structure serves laundering or sanctions evasion, and does not protect the central registry from a later breach — as the December intrusion of state hackers into the Treasury and FinCEN showed. The seal is an integrity and provenance layer for each declaration; the substantive truth of ownership still depends on registry verification, gatekeeper diligence, investigative journalism, cross-border data sharing and human judgment, none of which it replaces.

The promise and the peril

The promise is to give the registries back what is missing today: that every ownership declaration be **dated, attributable and court-proof**. When the next Pandora-scale leak erupts, no one will be able to pretend that what was declared was never declared; the journalist, the prosecutor and the authority will hold an unalterable receipt of who said what and when. In a world where a name on paper travels with no guarantee, making the declaration's origin unalterable protects the honest investigator and the citizen alike.

The peril calls for the usual nuances. First, do not confuse **integrity with truth**: a false, sealed declaration is still false — and now immutable besides; the seal could breed a false reassurance ("it is anchored, therefore it is true") that corresponds to nothing. Second, **privacy**: after the CJEU judgment, sealing declarations cannot become an excuse to expose personal data beyond legitimate interest; the fingerprint is anchored, not the data. Third, **concentration**: if only the big corporate service providers can seal solvenly, the standard could expel the small jurisdictions and actors instead of strengthening them. The proof of existence is a declaration governance layer, not a certificate of true ownership or a substitute for the diligence that knows how to audit a structure.

Moreover, it must be recognized that the same immutable infrastructure can become a weapon of corporate **lock-in** if mechanisms of openness and interoperability are not put in place. When the big platforms control the sealing protocols, any change or migration to a competing environment requires rewriting the whole history — a prohibitive cost that consolidates their dominance. Carissa Véliz warns that "today's tech giants may manage to lock in their position" [12], and this scenario materializes when regulators impose no portability rules or open formats. The risk is aggravated by the 2025 precedent, as we have seen, in which FinCEN exempted American companies from declaring beneficial ownership, emptying the registry and leaving the trust structure in the hands of a small number of private entities [13]. Without explicit guarantees against this chaining, registry transparency could regress, becoming a layer that reinforces concentration instead of lifting the smaller actors.

It's time to act

For commercial registries and authorities (FinCEN, AMLA, Companies House): require every beneficial ownership declaration to arrive with a sealed, dated fingerprint, so that backdating and silent manipulation are detectable. For corporate service providers and obliged entities (firms, notaries, banks): seal the declaration and its supporting evidence at the moment of filing, not afterward, and keep the .asice container as part of due diligence. For legislators and the FATF: recognize the sealed declaration as proof of the accuracy and currency Recommendation 24 already demands, and treat the certain date as the difference between a declaration and an alibi fabricated later. For investigative journalists and civil society (ICIJ, Open Ownership, Transparency International): be able to confront the current declaration with the version sealed on day zero, so that silent editing stops being invisible. And for honest declarants: embrace the proof of life and the forensic GPS as what distinguishes them from the front man who refuses to show his face.

As we saw in 2021, eleven million documents revealed what no registry could prove: who was really behind thirty thousand companies. Giving every ownership declaration a sealed receipt of who made it, what it said and when — never forgetting that the seal does not make what is declared in it true — is, on this front where the ownership of the world hides, time2act.

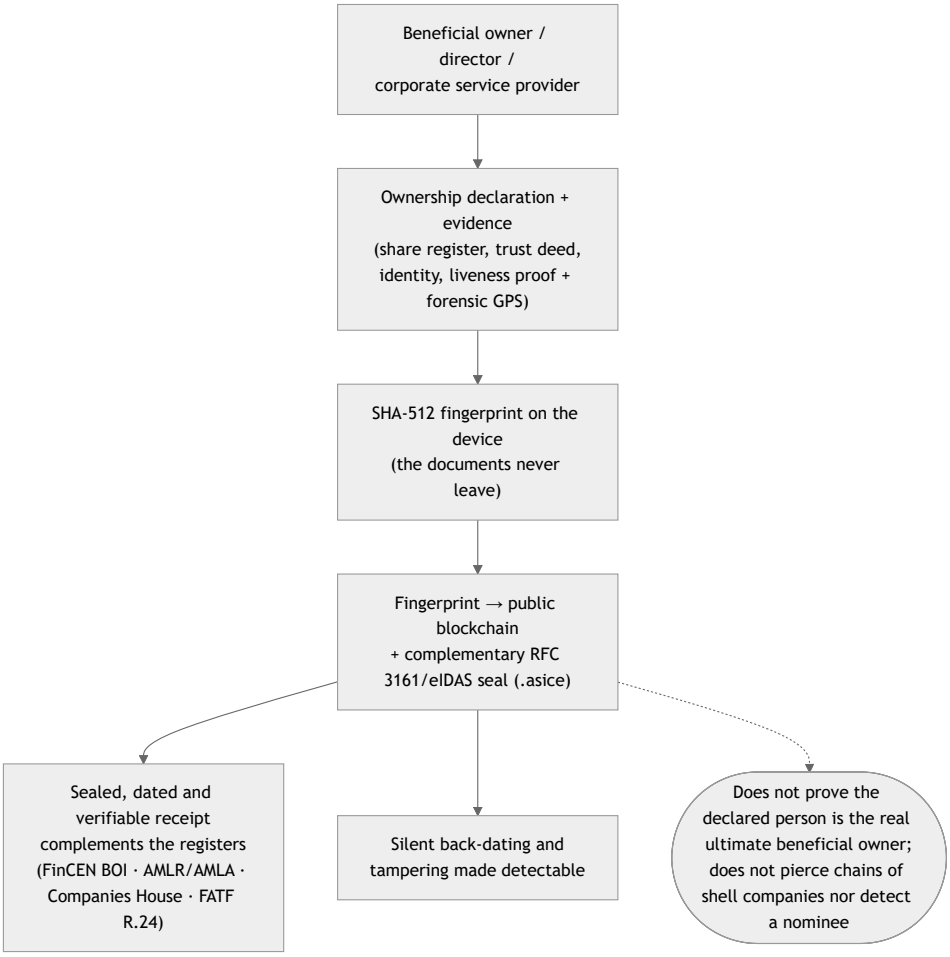


Figure 90.1 — The beneficial owner, the director or the corporate service provider files the declaration and its evidence (shareholder register, trust deed, identity, proof of life with forensic GPS in the field); the SHA-512 fingerprint is computed on the device — the documents do not leave — and anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). The receipt complements the registries (FinCEN BOI · AMLR/AMLA · Companies House · FATF R.24): it proves who declared what and when, and makes backdating detectable. Limit: the seal does not prove that the declared person is the ultimate beneficial owner; it does not pierce chains of shell companies or detect a front man. Sources: Pandora Papers (ICIJ); Panama Papers; Transparency International; FinCEN; CJEU C-37/20 and C-601/20; FATF R.24; Open Ownership.

---

[1] ICIJ — *About the Pandora Papers investigation* (October 3, 2021; 11.9 million documents, 2.94 TB, 14 providers, 29,000+ beneficial owners, 35 leaders and 330+ politicians). https://www.icij.org/investigations/pandora-papers/about-pandora-papers-investigation/ ↩

[2] Wikipedia — *Panama Papers* (April 3, 2016; 11.5 million documents from Mossack Fonseca, 214,488 entities, 150 investigations in 79 countries, shutdown in 2018). https://en.wikipedia.org/wiki/Panama_Papers ↩

[3] Transparency International — *Panama Papers: 10 years on* (~1,360M USD recovered; CTA reforms, EU registers, ROE; "the architecture of offshore secrecy remains intact"). https://www.transparency.org/en/news/panama-papers-10-years ↩

[4] FinCEN — *FinCEN Removes Beneficial Ownership Reporting Requirements for U.S. Companies* (interim rule of March 2025; exemption of domestic companies; covered population < 12,000). https://www.fincen.gov/news/news-releases/fincen-removes-beneficial-ownership-reporting-requirements-us-companies-and-us ↩

[5] Journal of Accountancy — *BOI update: 6.5 million of 32 million reports filed so far* (original estimate of 32 million obliged companies; 6.5M filed by November 2024). https://www.journalofaccountancy.com/news/2024/nov/boi-update-6-5-million-of-32-million-reports-filed-so-far/ ↩

[6] Transparency International — *EU court ruling on beneficial ownership registers: legitimate access* (CJEU judgment of November 22, 2022; invalidation of public access; the legitimate interest framework). https://www.transparency.org/en/blog/eu-court-ruling-on-beneficial-ownership-registers-legitimate-access ↩

[7] PwC Ireland — *The EU's new Anti-Money Laundering Authority* (AMLA Frankfurt operational since July 1, 2025; AMLR (EU) 2024/1624 applicable July 10, 2027; interoperable UBO registers). https://www.pwc.ie/services/audit-assurance/insights/eu-new-anti-money-laundering-authority.html ↩

[8] GOV.UK — *Companies House confirms identity verification rollout from 18 November 2025* (ECCTA; identity verification of directors and PSCs; 100,400 companies affected, 10,200+ filings rejected). https://www.gov.uk/government/news/companies-house-confirms-identity-verification-rollout-from-18-november-2025 ↩

[9] Transparency International US — *Unanimous Eleventh Circuit decision upholds the Corporate Transparency Act* (December 16, 2025; the highest court to uphold the CTA's constitutionality). https://us.transparency.org/news/unanimous-eleventh-circuit-decision-upholds-the-corporate-transparency-act/ ↩

[10] FATF — *Beneficial Ownership* (Recommendation 24 revised in 2022 and 2023 guidance; adequate, accurate and up-to-date information; a multi-pronged approach). https://www.fatf-gafi.org/en/topics/beneficial-ownership.html ↩

[11] Open Ownership — *Verification of beneficial ownership data: Overview* (self-declared data needs verification: compliance, cross-checks, supporting evidence). https://www.openownership.org/en/publications/verification-of-beneficial-ownership-data/overview/ ↩

[12] Carissa Véliz (ed.) — *The Oxford Handbook of Digital Ethics* (Oxford UP, 2022). ↩

[13] FinCEN (2025). https://www.fincen.gov/news/news-releases/fincen-removes-beneficial-ownership-reporting-requirements-us-companies-and-us ↩

Workforce Identity: Who We Really Hire Remotely

"Man is least himself when he talks in his own person. Give him a mask, and he will tell you the truth."

— Oscar Wilde, "The Critic as Artist" (1891)

In July 2024, **KnowBe4** — an American cybersecurity training company — hired a remote software engineer for its internal IT team without knowing that, on the other side of the screen, sat a North Korean operative. The candidate had used the **stolen identity of a real U.S. citizen** and a stock photo retouched with AI to match the application's headshot; with that he passed **four video-conference interviews**, a background check and reference checks. On July 15, hours after the corporate laptop reached his address, the company's **EDR** detected malware-loading attempts; the security operations center isolated the machine at 22.20 and, with Mandiant and the FBI, established that the supposed employee was a nation-state actor operating over VPN from North Korea or China, and that the laptop had ended up at a mule **"laptop farm"** on U.S. soil. No data was lost, but the chief executive, Stu Sjouwerman, published the incident on July 23 as a warning: hiring, and not just logging in, is now an identity attack surface [1].

This chapter is not about the citizen's digital identity — the foundational, governmental credential chapter 47 addresses. Here we speak of the **employment relationship**: the hiring, the interview, the identity verification, the onboarding and the continuity of remote work. The question is concrete: how do we prove that **the same verified person** was present at the interview, at the identity check and at the handover of credentials and laptop — and that whoever has access today is, still, that person? It is the book's thesis applied to the exact moment when a company decides to hand the house keys to a stranger.

Hiring is now an attack surface

The KnowBe4 case is not anecdotal; it is the pattern. On **July 24, 2025**, the U.S. Department of Justice sentenced **Christina Marie Chapman**, aged 50, of Litchfield Park, Arizona, to **102 months in prison** for running from her home a "laptop farm" that allowed North Korean IT workers to pose as U.S. citizens at **309 American companies** and two foreign ones, using **68 stolen identities** and generating more than **$17 million**. She hosted the corporate laptops in her house and shipped **49 abroad**, including to a Chinese city on the North Korean border [3].

The scale of the fraud keeps growing. In a coordinated DOJ action of **November 14, 2025**, it was established that facilitators based in the U.S. and Ukraine had helped DPRK IT workers obtain fraudulent remote employment at **more than 136 American companies**, compromising more than **18 real identities** and generating more than **$2.2 million** for the regime [4]. And it is no marginal phenomenon: the FBI, in alert I-012325-PSA of **January 23, 2025**, confirmed that North Korean workers use **AI and face-swapping technology** during video interviews to obfuscate identity, and that they have pivoted to **extorting** former employers, hijacking proprietary code and data [2].

The front door: the interview and the onboarding

What used to be a formality — seeing the candidate's face on screen — has become the blind spot. **Gartner** predicts that, by **2028**, one in four applications (25%) will be fake, driven by deepfakes, voice synthesis and chatbots; the survey indicates that **6%** of candidates admit to impersonation-style cheating in the interview [5]. Voice detection confirms it from the other side: **Pindrop's** data shows AI-generated or replayed audio in roughly **1 in 106** monitored meetings, with deepfake activity multiplied almost **13-fold (1,300% year-on-year)** and synthetic voices reaching **0.33%** of contact center calls in the fourth quarter, 173% more than in the first [7].

Recruitment leaders already live it as the top threat. According to a January 2026 study, **69% of UK hiring leaders** identify AI impersonation and deepfakes as the most sophisticated emerging threat to recruitment integrity (First Advantage); **Checkr** found that **23%** of companies detected identity fraud among new hires, and the **Greenhouse** report indicates that **91%** of U.S. hiring managers had encountered or suspected AI-generated interview answers [6]. The fraud no longer begins at login: it

begins at the first video call. The liveness handshake that armors the video call itself is described in chapter 88; here we deal with the whole hiring chain.

The regulatory framework

The institutional response is being built. As we have seen, in July 2025 NIST published the final version of the SP 800-63-4 Digital Identity Guidelines, which define the identity assurance levels (IAL1/2/3) and a new verification taxonomy — Remote Unattended, Remote Attended (for example, a video session), In-Person Unattended (kiosk) and In-Person Attended — a framework to which companies can map the rigor of onboarding identity verification [8]. The FBI, in the same alert and in its wanted and guidance pages on "fraudulent remote IT workers from the DPRK," recommends implementing identity verification during the interview, the onboarding and the entire remote employment relationship [2].

There is, moreover, an especially harsh layer of legal risk. OFAC sanctions against the DPRK make paying or employing a sanctioned North Korean worker a strict-liability violation: a hiring mistake becomes sanctions and money-laundering risk, with wire fraud, aggravated identity theft and IEEPA charges in the Chapman and Jin Sung-Il cases. Add to that U.S. work eligibility verification (Form I-9 / E-Verify), the legal gate the deepfake-retouched stolen identities aim to defeat, and in Europe the eIDAS / eIDAS 2.0 framework and the EUDI Wallet, with the transparency duties of the European AI Regulation.

And no minor limit: the GDPR and biometric laws like Illinois's BIPA restrict how much biometric or liveness data an employer can collect, which forces a privacy-preserving sealing.

Why the FBI's guidance alone does not save

It must be said bluntly. KnowBe4 had EDR, did video interviews, checked backgrounds and references — and it still hired the operative. The recommendations exist, but the problem is that **they leave no verifiable trace**: the video interview is held and forgotten; the identity verifier looks at a document and a face and draws a conclusion that lives, at best, in a note in an editable HR system; the laptop handover is a delivery slip. When, months later, the alarm goes off — "the person we hired is not the one who has access today" — there is no unalterable record to check against: there is HR's memory and a pile of screenshots. What is missing is a way to **bind, with a certain date and without storing sensitive data, the same verified person to the three critical moments** of the hiring.

The proof of existence applied

This is where **t2p** supplies the missing piece, and it does so on the hiring's identity event, not on a résumé claim. At three moments a capture is taken on the device: (1) at the live or video interview, a proof of life of the candidate's face plus the spoken consent statement; (2) at the identity verification, the identity document presented and the match against a selfie; (3) at the onboarding, a signed attestation binding that proven identity to the specific person who will hold the credentials and receive the laptop. Each capture is fingerprinted with **SHA-512 on the device** — the biometric or document file never leaves — receives an **RFC 3161 timestamp conforming to eIDAS**, is anchored to the **public blockchain** and packaged in an **.asice** container verifiable offline, with **forensic GPS** and the proof of life recording where and when the human really was. The unique value is **continuity and non-repudiation**: an unalterable, dated chain proving that the same verified human was present at the interview, at the check and at the device handover, so that a later dispute is tested against a sealed record and not against HR's memory.

This **complements, never replaces**, the existing controls: it coexists with NIST 800-63 IAL2/3 verification, I-9/E-Verify checks, background screening, the audits of duplicate résumés and reused phones or emails the FBI recommends, and real-time deepfake detection (Pindrop, Reality Defender, the C2PA/SynthID provenance **as we have seen**).

Now the limit, in all its rawness. "Garbage in, garbage on-chain": the seal proves that **a specific face, document and attestation** were captured at a place and instant by an author and have not been altered since; **it does not prove the person is honest**, or that the document is really theirs and not a high-quality stolen identity, or that they are not a sanctioned national, or that they will not exfiltrate data tomorrow or run a laptop farm. An impeccable deepfake or a real stolen identity that passes the proof of life produces a perfectly valid seal of a fraudulent act. The container guarantees the **integrity and**

provenance of the onboarding evidence; it replaces neither background screening, nor sanctions screening, nor deepfake detection, nor document authenticity verification, nor the continuous monitoring of insider threat and human judgment.

The promise and the peril

The promise is to turn remote hiring into an act with **verifiable memory**: being able to prove, in an incident like KnowBe4's, that the person at the interview, the one at the identity check and the one at the laptop handover were the same — or, precisely, to discover they were not — with proof that survives staff turnover and lost emails. In a market where a corporate laptop travels to a stranger without any identity receipt, making the chain unalterable protects the honest company and the legitimate candidate alike.

The peril calls for nuances. The first, do not confuse **integrity with truth**: an impeccable seal of a stolen identity is still the seal of a fraud, and it can breed a false reassurance — "it is anchored, therefore trustworthy" — that corresponds to nothing. The second, **privacy**: the face, the document and the proof of life are sensitive biometric data; the GDPR and BIPA demand that only the fingerprint leave, computed on the device, and never a central registry of candidates' faces. The third, **bias and exclusion**: the proof of life and document verification can fail more for certain groups or devices, and cannot become a barrier to employment. The proof of existence is a hiring non-repudiation layer, not a certificate of trust or a substitute for recruitment judgment.

Moreover, the danger of relying on a mere "click" as a form of consent widens when the information underpinning that agreement is so opaque and disproportionate that the user cannot really understand what they are authorizing. In the case of Facebook's facial recognition, Marc Goodman notes that the terms of use run to 9,300 words — a text so long it makes critical reading, and therefore genuine, verifiable consent, nearly impossible [11]. This information asymmetry creates a legal gap: the candidate can digitally sign an authorization without knowing which biometric traits are stored or how they may be reused. If the proof of existence does not include a mechanism that shows, unequivocally and auditably, that the individual has understood and accepted the conditions, the system is exposed to claims for GDPR and BIPA violations, and to an erosion of trust that outweighs any benefit of the immutable chain. Without verifiable, comprehensible consent, the technology becomes a new form of data exploitation rather than a guarantee of security.

One dimension the book *Reboot* puts on the table is the crisis of digital identity when platforms become markets of fictitious identities. The New York Attorney General's report (2021) revealed 9.3 million comments to the FCC from invented users, almost all created by a single 19-year-old, with close to 500,000 of Russian origin [12]. The solution Noveck proposes — a verifiable personhood proof generated entirely on the device — complements the neutral blockchain anchor: the immutable record ensures the proof has not been altered, while a "personal API" guarantees that only the device's owner can reveal their voice or digital presence, with no centralized biometric database. This closes the gap between the chain's immutability and the individual's autonomy, preventing fake identities from infiltrating critical processes and reinforcing real trust.

It's time to act

For **HR and recruitment teams**: seal the three moments — interview, identity verification and laptop handover — so the person's continuity has a certain date, without ever storing the biometric data. For **security leads (CISOs and SOCs)**: require the onboarding .asice as part of the employee file and check it against any insider threat alert, as KnowBe4 did with its EDR. For **identity verification and deepfake detection vendors** (NIST 800-63, Pindrop, Reality Defender): integrate on-device sealing with IAL2/3 verification and real-time detection, so proof and detection travel together. For **regulators and the FBI/IC3** (OFAC, eIDAS 2.0): recognize the sealed identity chain as evidence of due diligence in remote hiring and as proof against sanctions risk. And for **job platforms and remote talent marketplaces**: anchor the onboarding identity proof so that a substitution — the hired person is not the one who logs in — is detectable.

At KnowBe4, a nation-state operative passed four interviews and a background check before the laptop gave him away at 22.20; as we have seen, giving every remote hire a sealed proof binding the same verified person from the interview to the keyboard — never forgetting that the seal does not make the face in front of it honest — is, on this front where hiring is now attacking, **time2act**.

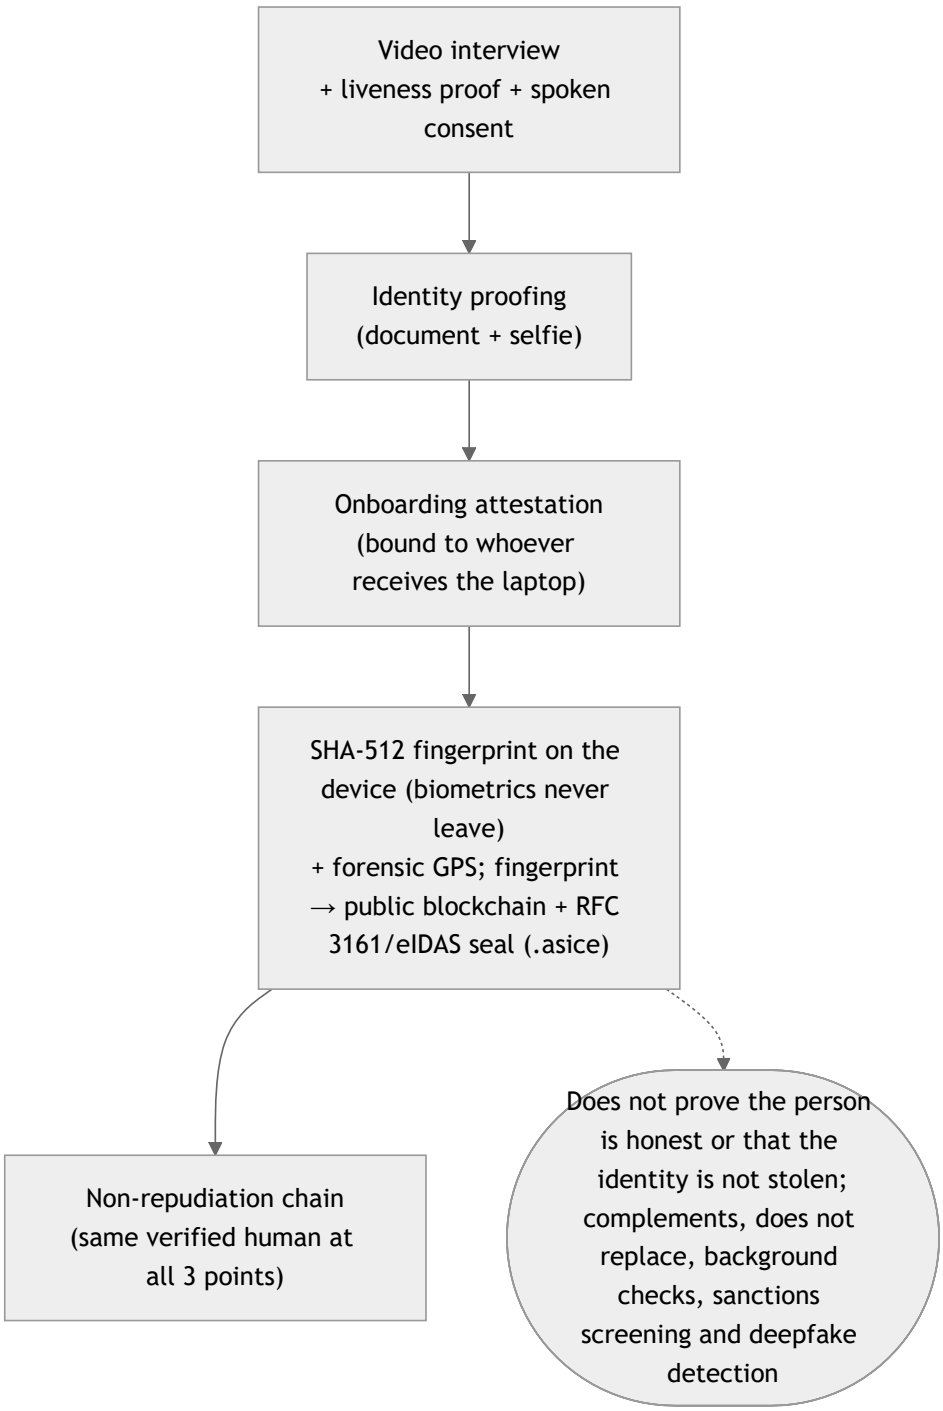


Figure 91.1 — Remote hiring sealed in three links: video interview with proof of life → identity verification (document + selfie) → onboarding attestation bound to whoever receives the laptop. At each moment the SHA-512 fingerprint is computed on the device (the biometrics do not leave), with an RFC 3161/eIDAS seal, forensic GPS and anchoring to the public blockchain (.asice). The result: a non-repudiation chain proving that the same verified human was at all three points. Limit: the seal does not prove that the person is honest, that the identity is not stolen or that they are not a sanctioned national; it complements, does not replace, background and sanctions screening and deepfake detection. Sources: the KnowBe4 case; the Chapman case (DOJ); FBI/IC3 I-012325-PSA; NIST SP 800-63-4; Pindrop.

---

[1] KnowBe4 — *How a North Korean Fake IT Worker Tried to Infiltrate Us* (Stu Sjouwerman; a first-person account: AI-retouched stock photo, four video interviews, EDR detection on July 15, 2024, mule laptop farm, FBI/Mandiant). https://blog.knowbe4.com/how-a-north-korean-fake-it-worker-tried-to-infiltrate-us ↩

[2] FBI/IC3 — *Public Service Announcement I-012325-PSA: North Korean IT Workers Conducting Data Extortion* (January 23, 2025; AI face-swapping in interviews, the pivot to extortion, identity verification recommendations). https://www.ic3.gov/PSA/2025/PSA250123 ↩

[3] US DOJ (USAO-DC) — *Arizona Woman Sentenced in $17M IT Worker Fraud Scheme for North Korea* (July 24, 2025; Christina Chapman, 102 months, 309 companies, 68 stolen identities, $17 million, laptop farm, 49 laptops shipped abroad). https://www.justice.gov/usao-dc/pr/arizona-woman-sentenced-17m-it-worker-fraud-scheme-illegally-generated-revenue-north ↩

[4] Crowell & Moring — *DOJ Announces Major Enforcement Actions Targeting North Korean Remote IT Worker Schemes* (November 14, 2025; more than 136 victim companies, more than 18 compromised identities, more than $2.2 million in proceeds). https://www.crowell.com/en/insights/client-alerts/doj-announces-major-enforcement-actions-targeting-north-korean-remote-it-worker-schemes ↩

[5] The Decoder — *Gartner predicts that by 2028, one in four job applicant profiles will be fake* (the 25% forecast for 2028; the 6% admission of interview impersonation, survey of 3,000 candidates). https://the-decoder.com/gartner-predicts-that-by-2028-one-in-four-job-applicant-profiles-will-be-fake/ ↩

[6] People Management — *Deepfakes and AI-enabled impersonation rank among top recruitment threats* (January 19, 2026; First Advantage 69% of UK leaders; Checkr 23% identity fraud; Greenhouse 91%). https://www.peoplemanagement.co.uk/article/1945557/deepfakes-ai-enabled-impersonation-rank-among-top-recruitment-threats-research-reveals ↩

[7] Pindrop — *The Growing Trend of Deepfakes in Interviews* (AI audio in ~1 in 106 meetings; ~1,300% year-on-year; synthetic voices in 0.33% of calls in Q4 2024). https://www.pindrop.com/article/growing-trend-of-deepfakes-in-interviews/ ↩

[8] NIST — *SP 800-63-4 Digital Identity Guidelines* (final, July 2025; IAL1/2/3 assurance levels and the Remote/In-Person × Attended/Unattended verification taxonomy). https://pages.nist.gov/800-63-4/sp800-63.html ↩

[9] FBI — *Fraudulent Remote IT Workers from DPRK* (wanted/guidance; use of stolen identities, intermediary computers and AI, with mitigations for employers). https://www.fbi.gov/wanted/cyber/fraudulent-remote-it-workers-from-dprk

[10] Microsoft Security — *Fake Employees, Real Threat: Decentralized Identity to combat Deepfake Hiring* (verifiable credentials and continuous verification of onboarding identity; adjacent-solution context). https://techcommunity.microsoft.com/blog/microsoft-security-blog/fake-employees-real-threat-decentralized-identity-to-combat-deepfake-hiring/4479864

[11] Marc Goodman — *Delitos del futuro* (Ariel, 2015). ↩

[12] New York Attorney General (2021). https://ag.ny.gov/press-release/2021/attorney-general-james-issues-report-detailing-millions-fake-comments-revealing ↩

Platform Work: When the Boss Is an Algorithm

"...a search, too, for daily meaning as well as daily bread, for recognition as well as cash... for a sort of life rather than a Monday through Friday sort of dying."

— Studs Terkel, Working (1974)

You open the app first thing to start your shift and the screen shows a single line: "Account deactivated." No name, no phone number, no button to ask for an explanation. The decision that costs you your wage was made by an automated system, out of your sight. This is the starting point of the case known as "Robo-Firing" — Drivers v. Uber and Ola — in which a group of drivers had to litigate for years, up to the Amsterdam Court of Appeal (2023), just to learn why: the court, invoking Article 22 of the General Data Protection Regulation, upheld their appeal in three linked judgments and ordered Ola to explain the logic of its fully automated decision, and Uber to give meaningful information about the algorithmic dismissals; Uber ended up ordered to pay €584,000 in coercive fines — plus €4,000 for each day of noncompliance — for failing to comply with the order to explain itself [4]. Think about that for a moment: it took a European appeals court and daily fines for a worker to learn the reason for his own dismissal.

This chapter deals with work managed by digital platforms — couriers, drivers, microworkers — and the employment relationship with a boss that is an algorithm. Who the person we hire remotely really is, is dealt with in chapter 91; labor conditions in the global supply chain, in chapter 102; what is inside AI models, in chapter 79. The key question: when the boss is an algorithm, who takes the minutes of the employment relationship? The answer: both parties, with anchored evidence — the worker's sealed work diary and the platform's auditable decision log, with **t2p**: the SHA-512 fingerprint on the device, the RFC 3161/eIDAS timestamp and the anchor to the public blockchain inside an .asice container [5] — so that the rights Europe and Spain have already written on paper [1][2] can be proven in fact.

The boss who does not show his face

Algorithmic management deserves to be described without demonizing it. The algorithm assigns jobs in real time, sets dynamic rates by supply and demand, scores performance, orders the work queue and, at the limit, disconnects. It is efficient — that is why it exists: without this machinery, no delivery or ride app could coordinate millions of users. But the efficiency has a hidden cost: a radical informational asymmetry.

The platform records everything — positions, response times, acceptances, clicks — in massive logs that grow every second. The worker, by contrast, sees only the day's interface: the jobs that appear, their score, that moment's rate. And tomorrow the configuration may have changed without notice. When the conflict comes — before the Labor Inspectorate, before a judge — one party appears with the system's complete log and the other with memories.

The diagnosis is already law in Europe. Directive (EU) 2024/2831, in force since December 2, 2024 and with transposition due by December 2, 2026, imposes transparency on the monitoring systems and the algorithms that affect working conditions; requires human oversight of the decisions and a right to human intervention in the important ones; and forbids using sensitive data — emotional state, biometrics, private conversations — in automated decisions [1]. And it establishes the keystone: the presumption of employment when facts indicating direction and control are found, with the burden of rebutting it on the platform [1].

Note the decisive word: "facts." The Directive is a machine for turning facts into rights. The problem is that, in this sector, the facts are in the exclusive custody of one of the parties. This is exactly the hole the proof of existence closes.

Spain as laboratory: the Rider Law and the Glovo case

Spain got there first. Law 12/2021 — the "Rider Law," in force since August 2021 — established the presumption of employment for delivery work when the company exercises its organization, direction and control through the algorithmic management of the service or of working conditions, and created a pioneering right: the works council must be informed of the parameters, rules and instructions of the algorithms that affect working conditions [2].

And the case that put the law to the test was Glovo. The Barcelona courts opened criminal proceedings against the founder and chief executive, Óscar Pierre, over the false self-employment model; the claims for unpaid contributions and the Labor Inspectorate's fines added up to some €200 million, and the competitor Just Eat claimed 295 more for unfair competition [3]. The outcome came between late 2024 and 2025: Glovo announced the move to the employment model — between 10,000 and 15,000 couriers, combining direct hiring and fleets — and Pierre acknowledged before the judge that the regularization sought "social peace" and to avoid "arguments" with the Inspectorate, and that the criminal and economic pressure had forced the change: "the cash box doesn't let you operate." The parent company Delivery Hero raised its provisions for Glovo Spain to between €440 and 770 million and estimated some €100 million of impact on the 2025 operating result [3].

The evidentiary lesson — which is this book's, not a political reproach — is this: it took years of inspections, reports and data cross-checks to reconstruct, from outside and after the fact, an employment reality that happened every day inside an app. Every euro of fine is also the price of evidence that took years to dig out. With records anchored from day one, the same truth would have been visible within weeks — for all parties, including the platform that might have wanted to prove its diligence.

Dismissal by algorithm

Let us return to "robo-firing," because it condenses the inversion of labor grammar. In the classic world, the employer dismisses and must give reasons; in the platform world, the system disconnects and it is the worker who must litigate — for years, if need be [4] — just to obtain the reasons. The European Directive responds by generalizing what Amsterdam had to impose case by case: human oversight and a right to human intervention in the important decisions [1].

But even with the reasons in hand, the second step remains: testing them. If the algorithm says "you rejected too many orders" or "customers rated you poorly," what does the worker rebut it with, having no records of his own? The platform holds the data and the criterion; he, not even his own version of the facts.

The answer has two symmetrical halves — and it matters that it is both. The worker needs his indelible work diary: his version of the facts, sealed day by day on his device [5]. And the platform needs — by legal obligation [1][4] and by self-interest — an auditable decision log: which rule was applied, on what data, with what human review and when. Both halves are the same technology; neither can be rewritten when the conflict comes.

The proof of existence applied

The same shift, recorded from both sides of the screen — and neither version can be rewritten.

1. The sealed work diary

Each working day, the worker seals his day on his device: hours connected, availability, jobs offered, accepted and rejected, instructions received, captures of the conditions and rates the app showed that day [5]. It is not spying: it is his own working life, documented from his side of the screen. The result is exactly what the presumption of employment needs — dated facts indicating direction and control [1][2] — immune to the argument "the app didn't say that": that morning's capture carries a fingerprint and a certain date. And if one day the disconnection comes, his chronology does not start from zero: it already exists, anchored, from the first shift.

2. The auditable deactivation

The platform that so decides — or that the law obliges [1][4] — can anchor its decision log: when the algorithm flagged the account, what data and what rule it applied, which human reviewed it and what they resolved. Each step, fingerprinted and dated [5]. Before the Amsterdam judge, that chain would have been the answer that cost €584,000 not to have [4]; against an accusation of discrimination or arbitrariness, it is the diligent platform's defense; and before the worker, it is the right to know — Article 22 of the GDPR, the Directive [1] — fulfilled by design, not by judgment.

3. The inspection that does not start from zero

The Glovo case took years of reports and reconstructions to establish an employment relationship the app knew minute by minute [3]. With workers' sealed work diaries and platforms' anchored decision logs, the Labor Inspectorate and the courts work on verifiable chronologies from the first request: fewer years of investigation, fewer reconstruction expert reports, and sanctions — or acquittals — resting on facts with a certain date [5]. The presumption of employment [1][2] stops being an evidentiary battle and becomes a check.

4. The algorithm with minutes

The Rider Law gives the works council the right to know the parameters, rules and instructions of the algorithms [2]; the Directive imposes continuous transparency [1]. But a document handed over on a Tuesday says nothing about Friday's silent change. The solution is sealed versioning: every version of the parameters handed over, fingerprinted and dated [5], so that any later modification is visible and datable. It is the same provenance logic the book applies to AI models in chapter 79, brought to the employment relationship: the algorithm can change, but it cannot make anyone believe it was always the way it is now.

The limits, stated clearly: the seal does not decide whether a relationship is employment — laws and judges do that [1][2]; it does not by itself open the algorithm's black box — the transparency duties do that [1][2] and, when needed, the courts [4]; and it replaces neither the collective agreement, nor the union, nor collective bargaining. The proof of existence makes the right enforceable; the right, people must make.

The promise and the peril

The promise: workers who no longer argue with a mute screen [4]; platforms that can prove their diligence instead of paying for it by judgment [3]; inspections that take weeks and not years; and a presumption of employment [1][2] that switches on with dated facts, not battles of witnesses. Algorithmic management, with minutes.

The perils, with the same clarity.

The first is **inverted surveillance**. It would be a tragic contradiction if the tool meant to balance the relationship ended up justifying more monitoring of the worker. The Directive already bans sensitive data — emotions, biometrics, private conversations — in automated decisions [1]; the sealed work diary must be the worker's and for the worker: on his device, under his control, with voluntary disclosure and revocability [5]. No platform may demand it or access it.

The second is **the biased record**. Whoever seals chooses what they seal: a platform could anchor only what favors it, and so could a worker. The seal attests to integrity and date, not completeness. That is why the design must be bilateral, and regulators must be able to demand the complete logs from the platforms — remember that the gaps get dated too [5].

The third is **portable reputation**. If the sealed work diary degenerates into a "worker score" exportable between platforms, we will have built the perfect blacklist — immutable, no less. The chronology is for defending rights in a conflict, not for serving as a permanent conduct record: expiry, revocability and an express ban on using it as a hiring filter.

And the fourth is **labor techno-solutionism**. No fingerprint raises a rate or pays a contribution. Platform work rights were born of the law [1][2], of the Inspectorate [3], of the courts [4] and of collective organization; cryptography only guarantees that, when proof is needed, it exists. Confusing the tool with the struggle would repeat, in a new sector, the oldest mistake.

It's time to act

If the algorithm directs, let there be minutes — of its orders and of your work.

For platform workers: seal your work diary — hours, jobs, instructions, captures of conditions [5]; know the presumption that protects you [1][2]; and if the disconnection comes, demand the reasons Europe already grants you [1][4].

For the platforms: anchor the decision log and the human review [1][4][5]; the Glovo case teaches that reality surfaces in the end, and that doing it late costs hundreds of millions [3] — being auditable from day one is cheaper than being made auditable by judgment.

For unions and works councils: demand the algorithms' parameters [2] with sealed versioning [5] — transparency without a certain date is an old photograph.

For the Labor Inspectorate and the judges: request and weigh the anchored chronologies of both parties [3][5]; let the investigation start with dated facts, not terabytes to dig through.

For legislators: when transposing the Directive before December 2, 2026 [1], build in the duty of an anchored decision log and the worker's right to his own sealed diary — and armor it against any use as a reputational filter.

It took the Amsterdam drivers an appeals court and €584,000 in fines for an app to tell them why it had fired them [4]. Let the next disconnected shift have the answer waiting: dated, anchored and enforceable. For decent work in the age of the algorithm, **time2act**.

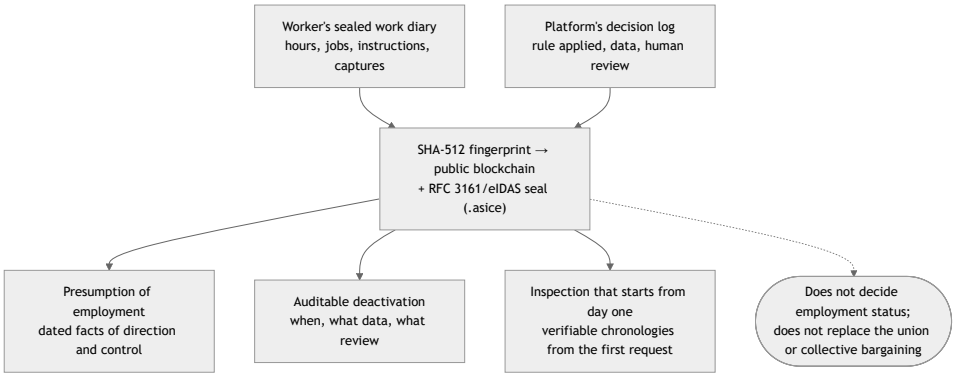


Figure 92.1 — The platform employment relationship with proof of existence: the worker seals his work diary (hours, jobs, instructions, captures of conditions) on his device; the platform anchors the log of algorithmic decisions and the human review (Art. 22 GDPR; Directive (EU) 2024/2831); the works council receives the algorithm's parameters with sealed versioning (the Rider Law). Three outcomes: the presumption of employment proven with dated facts, the auditable deactivation and the inspection that does not start from zero. Limits: the seal does not decide employment status or replace collective bargaining; the diary is the worker's, never a reputational filter. Sources: OJEU (Dir. 2024/2831); BOE (Law 12/2021); elDiario.es (the Glovo case); Worker Info Exchange (Drivers v. Uber/Ola).

---

[1] OJEU / BOE — *Directive (EU) 2024/2831 of the European Parliament and of the Council, of October 23, 2024, on improving working conditions in platform work* (in force December 2, 2024; transposition by December 2, 2026; presumption of employment on facts of direction and control with the burden of proof on the platform; ban on sensitive data in automated decisions; transparency, human oversight and human intervention in important decisions); Laboral Social — *Directiva sobre plataformas digitales de trabajo*. https://www.boe.es/buscar/doc.php?id=DOUE-L-2024-81667 · https://eur-lex.europa.eu/eli/dir/2024/2831/oj?eliuri=eli:dir:2024:2831:oj&locale=es · https://www.laboral-social.com/directiva-plataformas-digitales-trabajo-mejora-condiciones-laborales ↩

[2] BOE — *Law 12/2021, of September 28* (the "Rider Law": presumption of employment in delivery via algorithmic management; the works council's right to be informed of the parameters, rules and instructions of the algorithms affecting working conditions; in force since August 2021); PwC Periscopio — *¿Ha cumplido sus objetivos la Ley Rider?*. https://www.boe.es/buscar/act.php?id=BOE-A-2021-15767 · https://periscopiofiscalylegal.pwc.es/luces-y-sombras-de-la-ley-rider/ ↩

[3] elDiario.es — *El jefe de Glovo admite a la jueza que la regularización de los 'riders' busca evitar «discusiones» con Inspección de Trabajo* (criminal proceedings in Barcelona; ≈€200 million in contributions and fines; Just Eat claims €295 million); elDiario.es — *El jefe de Glovo dice que la presión penal y las multas forzaron cambiar su modelo laboral: «La caja no te da para operar»* (10,000-15,000 riders; Delivery Hero provisions of €440-770 million; ≈€100 million impact on 2025 EBITDA); Crónica Global — *Óscar Pierre justifica el cambio de modelo laboral por «presión económica y penal»*. https://www.eldiario.es/catalunya/jefe-glovo-admite-jueza-regularizacion-riders-busca-evitar-discusiones-inspeccion-trabajo_1_11870197.html · https://www.eldiario.es/economia/jefe-glovo-dice-presion-penal-multas-forzaron-cambiar-modelo-laboral-caja-no-da-operar_1_12062602.html · https://cronicaglobal.elespanol.com/business/20250218/oscar-pierre-de-glovo-justifica-laboral-economica/925157487_0.html ↩

[4] Fountain Court Chambers — *Amsterdam Court Upholds Appeal in Algorithmic Decision-Making Test Case: Drivers v Uber and Ola* (April 2023; three judgments; Ola must explain the logic of the fully automated decision, Art. 22 GDPR); Worker Info Exchange — *Uber ordered to pay €584,000 for failure to comply with court order in robo-firing case* (plus €4,000/day); UK Human Rights Blog — *Amsterdam Court orders reinstatement of Uber drivers dismissed by algorithm*. https://fountaincourt.uk/2023/04/amsterdam-court-upholds-appeal-in-algorithmic-decision-making-test-case-drivers-v-uber-and-ola/ · https://www.workerinfoexchange.org/post/uber-ordered-to-pay-584-000-for-failure-to-comply-with-court-order-in-robo-firing-case/ · https://ukhumanrightsblog.com/2021/05/18/amsterdam-court-orders-reinstatement-of-uber-drivers-dismissed-by-algorithm/ ↩

[5] t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; content on the device, voluntary disclosure; the series' gaps get

dated). See chapters 13 and 15; the provenance of AI models, in chapter 79. ↩

Working Hours and the Labor Dispute: Hours on the Record

"Power concedes nothing without a demand. It never did and it never will."

— Frederick Douglass, Canandaigua speech (1857)

19:40 in an ordinary industrial park. The last order went out late, and the time record — the paper sheet in the binder, or the company's app — reads 18:00 sharp, like every day. In the cafeteria of a shopping mall, a waitress mentally adds up the hours she will not be paid this month. Neither gesture will make any noise today; both will weigh on the day there is a dispute.

The data says these are not anecdotes [1]: according to the Labor Force Survey, more than 40% of the overtime hours worked in Spain go unpaid. And that is with a working-time record that has been mandatory for all companies since 2019. The problem is not that the record does not exist: it is who holds it. Paper takes whatever is written on it; the company's digital system is controlled by the company. When the conflict erupts — the claim for hours, the change of conditions communicated by word of mouth, the dismissal that arrives right after a complaint — the labor court is left with two memories and one party's record.

This chapter deals with the most basic variable of the country's most common contract — and the worst proven. Work governed by an algorithmic boss — couriers, drivers — is in chapter 92; the identity of who you hire remotely, in chapter 91; labor conditions in the global supply chain, in chapter 102; harassment and the technique of recording one's own conversations, in chapter 75. The key question: how do you argue about working hours when the official clock belongs to one of the parties? The answer: with a certain date on both sides — the worker's work diary and the neutral certification of the official record — through **t2p**: hash, time and anchor — the tripod of chapter 13 [3]. And with a context that pushes: the coming reform of the digital time record will demand exactly this architecture.

The record that is coming: immutable by law

The legislative news [1]: on September 30, 2025 the government approved the urgent processing of a decree for the digital working-time record — decoupled from the reduction of the working week — with the aim of passing it during 2026 (no entry-into-force date published yet). What the draft requires, according to the public documentation: traceable, immutable electronic records; timestamping; recording of breaks; remote, real-time access for the Labor Inspectorate; and a history of modifications with author and date — a unilateral correction without the worker's knowledge may be treated as manipulation before a court [1].

Pause on what this means for this book's thesis — and set it beside VeriFactu, the system chapter 34 explains: for the second time in a few years, the state is demanding by rule immutability, traceability and a certain date to solve a problem of trust. First invoices; now hours. The proof of existence is no longer a proposal: it is the direction of travel.

And now, what the reform does not give. The record, however immutable, lives in the systems of the company or its provider; the remote access is for the Inspectorate, not for the worker, who still has no opposable copy of his own hours. And the honest company, facing an accusation of manipulation, must prove a negative — that it touched nothing — with its own systems as the only witness. The missing layer looks both ways: the copy with a certain date in the worker's hands, and the neutral certification of the official record that makes the company's compliance provable, not merely arguable.

The proof of existence applied

The labor dispute is won or lost on dated facts; let each side have its own [3].

1. The worker's work diary

On his personal device, the worker seals what affects him: the photo of today's clock-in, the note of the hours worked, the published shift roster [3]. Thirty seconds at the end of the day. Only the fingerprints are anchored; the content stays on his phone, and he need show it to no one — until the day it is needed,

before the union delegate, the Inspectorate or the labor court. Months of entries with a certain date weigh more than any memory — and than any one-sided record.

And a mention this chapter cannot omit: nowhere is this diary worth more than in domestic work. Household workers — the most informal contract in the country: hours agreed by word of mouth, cash, often no record of any kind — have, in the same thirty-second gesture, the first labor evidence of their history [3]. The dignity of a job is also measured by its capacity to be proven.

2. The clock-in with place and time

When the official system "fails" — the day that curiously does not appear, the clock-in someone corrected — the entry and exit photo with t2p's forensic GPS — a genuine fix, accuracy of about 50 meters, maximum age of about 30 seconds; a simulated location blocks certification [3] — attests to where the worker was and when. The same goes for the unpaid commute: the real workplace, dated. And the red line, stated right here: the moment of clocking, never continuous tracking — the working day is documented, the person is not surveilled.

3. The roster and the order

The published roster, captured with a certain date; the roster changed at the last minute, likewise; the order received by message or email, certified [3]. The "I never told you that" and the change of conditions communicated by word of mouth collide with dated screens. And if the conversation is face to face, recording one's own conversation is lawful — the technique and the limits are in chapter 75.

4. The certified official record — the company's symmetry

The company that does things right closes its time record every month with a batch certification: the monthly file's fingerprint, chained to the previous one [3] — the same logic VeriFactu imposes on invoices. Before the Inspectorate — which will be able to access it remotely [1] — and before the court, its record is not merely immutable from now on: it is provably unmanipulated since each closing. Against the claim of phantom hours, against the accusation of retroactive correction, the honest company has dates, not pleadings. And the history of modifications the reform requires [1], anchored, becomes its permanent alibi.

The limits, in black and white: the seal attests to hours and communications — it does not decide whether the hour was overtime, complementary or voluntary: the collective agreement and the judge say that. The worker's personal diary does not replace the company's legal record — it complements and checks it. And proportionality rules: neither party may turn the working day into surveillance — no continuous geolocation of the worker, no ambient recording of the workforce.

The promise and the peril

The promise: labor trials that argue law, not facts [1][3]; the 40% of unpaid overtime [1] with a claiming tool within reach of any pocket; compliant companies that pass inspections with provably intact records; and the culture of "I'll write it down later" replaced by that of the dated fact — on both sides of the payslip.

The perils, without detours.

The first is **surveillance disguised as a record**. The step from the dated clock-in to the tracked worker is short and tempting. The legal and moral line: the moment is recorded — entry, exit, break — never the route; the record serves the working day, not the control of a life.

The second is **institutionalized distrust**. If every order must be captured and every shift sealed, the employment relationship is judicialized in advance. The answer is this book's usual one: the trail is for the conflict, not for the day-to-day — and the more verifiable the official record is, the less the parallel one is needed.

The third is **the gap of the most precarious**. Those who suffer unpaid hours most are those with the fewest tools and the least margin to document them without fear. The personal diary helps; but the real protection is the Inspectorate with remote access [1], the union delegates and the collective agreements — individual proof does not replace collective protection.

And the fourth: **confusing measuring with solving**. The perfect record does not pay the overtime: it makes it visible and claimable. If claiming still costs the job, the problem is not evidentiary. What is

needed is inspection, sanctions and a culture where paying the hours worked is the minimum, not a conquest.

It's time to act

Let the clock of work have two owners.

For you, the worker: thirty seconds at the end of each day — the clock-in, the hours, the shift [3]; the day you need to claim, you will have months of dated facts. And if the matter is harassment, chapter 75 is yours.

For companies: certify the time record in chained batches [3] before the reform [1] half-asks it of you — the provably intact record is your best defense before the Inspectorate, the courts and unfounded claims.

For unions and delegates: make the sealed work diary a basic tool of membership and defense [3]; and negotiate into the agreements the worker's right to a verifiable copy of his record.

For the Labor Inspectorate: the remote access the reform will give you [1] is worth as much as the record's integrity — demand it anchored, with a verifiable history of modifications.

For legislators: the digital record reform [1] is the chance to close the circle: an immutable record, yes — and a copy with a certain date for the worker, and recognized evidentiary value for the neutral certification. Let the second half of the proof not be left in the inkwell.

Let the 19:40 of the industrial park and the cafeteria's hours be on the record — and let the next labor trial argue justice, not memory. For the hours that are life, **time2act**.

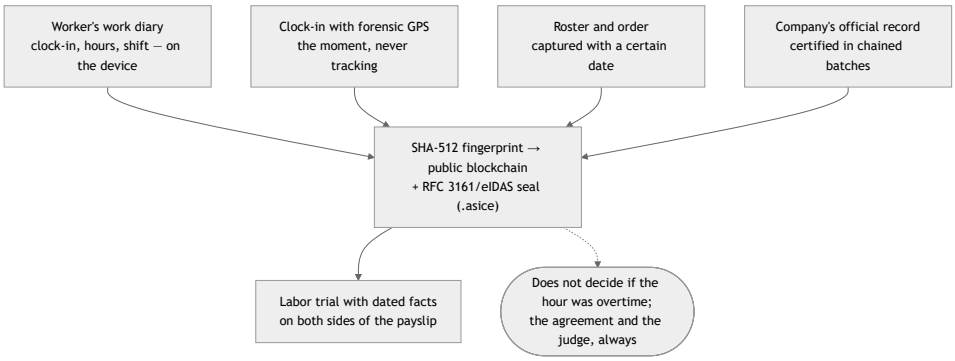


Figure 93.1 — The working day with proof of existence (sealed with t2p — chapter 13): the worker's work diary (clock-ins, hours, shifts — on his device, only fingerprints on the chain), the clock-in with forensic GPS (the moment, never tracking), the roster and the order captured with a certain date, and the company's official record certified in chained batches — provably intact before the Inspectorate and the court. Context: the digital time record reform (urgent processing since 30-9-2025) will demand immutability, traceability and remote access for the Inspectorate; >40% of overtime goes unpaid (EPA). Limits: the agreement and the judge decide the substance; proportionality always. Sources: EPA/INE; the digital time record draft.

---

[1] Ministry of Labor / trade press — *Registro horario digital obligatorio: la reforma en trámite* (urgent processing approved by the Council of Ministers on 30-9-2025, decoupled from the working-week reduction; aim of approval during 2026, no entry-into-force date published; the draft's requirements: traceable, immutable electronic records, timestamping, recording of breaks, remote real-time access for the Labor Inspectorate, a history of modifications with author and date); EPA/INE — more than 40% of overtime goes unpaid; RD-ley 8/2019 — mandatory daily working-time record. https://www.protime.eu/es-es/noticias/registro-horario-digital-obligatorio-2026 · https://factorial.es/blog/nuevo-registro-horario-digital-espana/ · https://www.inwout.com/post/registro-horario-digital-obligatorio-2026 · https://kronjop.com/es/newsroom/control-horario/digital/obligatorio/ ↩

[2] The general labor framework (the labor court claim; the Labor Inspectorate's reports; the delegates' role). See chapter 92 for platform work and chapter 102 for the labor supply chain.

[3] t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; certified screenshot and email; certify any file — only the fingerprint travels; background certification; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; one's own conversations and harassment, in chapter 75. ↩

Construction and Infrastructure: Who Proves How It Was Really Built

"All these must be built with due reference to durability, convenience, and beauty."

— Vitruvius, De architectura, I, 3 (first century BC)

On February 6, 2023, two earthquakes (Mw 7.8 and 7.5) shook southeastern **Turkey** before dawn. More than **55,000 people** died and more than **160,000 buildings** collapsed, many of them structures that had been legalized through repeated zoning amnesty programs (*imar barışı*) and raised with dubious, uncertified concrete. The rawest case is the **Isias Grand Hotel**, seven stories in **Adıyaman** that folded like an accordion in a few seconds and killed **72 people**, among them a volleyball delegation from a Turkish Cypriot school: 35 of the dead were children, teachers and coaches. In December 2024, after a seventeen-hour hearing, a court sentenced the owner **Ahmet Bozkurt**, his son **Mehmet Fatih Bozkurt** and the architect **Erdem Yılmaz** to up to **18 years** each for "causing death by conscious negligence." A 150-page technical report concluded that the hotel had been built with poor concrete mixed with river sand and gravel, had no ground study report at all, and had had a floor added without permission in 2016 [1].

This chapter is not about who owns a property or its legal title — that is chapter 35, the land registry. Here we speak of the **record of physical integrity**: the material tests, the site inspections, the geotechnical studies and the as-built drawings — that is, the proof of **how something was really built and how it was verified**, not of who owns it. The question is uncomfortable because a finished building hides its own history: once the concrete covers the rebar, no one sees whether the bars were there, whether the concrete had the declared strength, or whether the laboratory test was done honestly. And when this history is falsified, money is not what is lost: lives are.

The history buried in the concrete

The Turkish case is no exotic anomaly. An estimated **13 million Turkish buildings** — more than half the stock — were raised in breach of the rules, and at least **75,000 structures** in the affected area had received zoning amnesties before the earthquake [1]. These amnesties — the *Yapı Kayıt Belgesi*, registration certificates that retroactively legalized buildings in violation of the seismic codes — were described by Turkey's engineering associations as "an invitation to death." The underlying problem is that the paper that was supposed to prove a building safe either did not exist or could be bought after the fact.

The same pattern appears far from any seismic fault. On June 24, 2021, in **Surfside**, Florida, the **Champlain Towers South** tower collapsed in the early hours and killed **98 people**. The investigation by **NIST** — the U.S. federal agency — placed the start of the collapse in a **punching shear failure** of the pool deck slab: the concrete was weaker than the code required, and the rebar had been placed with excessive cover that reduced its shear capacity [3]. One slab, one construction detail, years of undocumented degradation: 98 dead. The difference between Adıyaman and Surfside is not the physics; it is that in both cases the **record of what was actually built** was nonexistent or could not be trusted.

When the material test is the lie

There is an even more direct fraud: falsifying the very test that is supposed to guarantee the material. The most spectacular case is the **Hong Kong–Zhuhai–Macao Bridge**, one of the most expensive works on the planet. Between 2012 and 2017, the **ICAC** (Hong Kong's anti-corruption commission) charged **19 workers** of the laboratory Jacobs China Limited with falsifying **384 concrete cube test reports**: they manipulated the computer's date and substituted weak cubes with high-strength ones, and even with a metal calibration cylinder. The cleanup cost taxpayers **HK$58 million**, and in June 2024 the director Albert Leung Wing-keung was jailed for two and a half years for defrauding the government of some HK$2 million [4].

In the United States, the president of **Testwell Laboratories**, V. Reddy Kancharla, was convicted of faking concrete strength tests on **more than 100 projects** in New York, among them the **Freedom Tower**, **Yankee Stadium** and the Second Avenue subway. When real cores were extracted and tested, the Freedom Tower concrete that had been certified at 12,000 PSI turned out to be, in fact, around 9,000

PSI. Kancharla received seven to 21 years in prison and $225,000 in restitution [5]. The technique was simple and devastating: instead of testing the cube and writing down the result, the specification value was written directly; the document said what it was supposed to say, but the physical reality was another.

And when it is not the material, it is the inspection. On August 14, 2018, the **Morandi Bridge** in **Genoa** collapsed and killed **43 people**. In July 2022, **59 defendants** went to trial; the prosecution concluded that some inspection documents were "clearly false" and that the operator, Autostrade, inspected itself through its subsidiary Spea. As we have seen, on April 7, 2023 a judge accepted a €30 million settlement between Autostrade and Spea [6]. Self-surveillance — being judge and party of one's own safety — is the hole the falsification slips through.

The law already demands the thread; the anchor is missing

Regulation has reacted, and forcefully. In the United Kingdom, after the **Grenfell Tower** fire — as we have seen — the **Building Safety Act 2022** requires maintaining a digital "golden thread": a traceable, version-controlled record of the design, the products, the materials, the approvals and the as-built information, kept for the whole life of higher-risk buildings [7]. In the European Union, **Regulation (EU) 2024/3110** — the revision of the Construction Products Regulation, in force since January 7, 2025 and applicable to most products from January 8, 2026 — makes the **Digital Product Passport** a mandatory legal requirement for CE marking: it must carry the Declaration of Performance and Conformity and the technical documentation through a QR code or data carrier [8]. Add the **ISO 19650** series, the international standard for managing information with **BIM** across the whole life cycle and producing an as-built "digital twin" of the asset. The law, then, already asks for exactly what is missing: a record of *how it was built*. But a golden thread stored on a server controlled by the party who must account for it is, in a dispute, one side's paper.

Why the golden thread alone does not save

It must be said bluntly. A digital record, however well designed, is only as reliable as the impossibility of silently altering it. If the pre-pour inspection report lives in a database the operator can edit — as Autostrade inspected itself — tomorrow the date can be touched up, a photo replaced or an awkward test deleted without a trace. The Hong Kong fraud consisted, literally, of **changing the computer's date**; Testwell's, of **writing the specification value** instead of the measured one. No "golden thread" that can be rewritten after the fact stops these two techniques. What the law demands is a record; what is missing is for that record to be **unalterable, independently dated and bound to an author and a place**, so that any later substitution or backdating can be proven.

The proof of existence applied

This is where **t2p** supplies the missing layer at every irreversible milestone of the works. **First**, before the pour, the field inspector photographs or films the rebar, the formwork and the cover; the device computes the **SHA-512 fingerprint** on the fly, obtains an **RFC 3161 timestamp conforming to eIDAS**, anchors it to the **public blockchain** and adds the **forensic GPS** and a **proof of life** binding the inspector to that place and instant. **Second**, the materials laboratory seals every cube test result — and the instrument's raw data — at the exact moment of the 28-day break, so the date and the value can no longer be backdated or swapped (precisely the techniques of Hong Kong and Testwell). **Third**, the engineer seals the signed as-built drawing set and the geotechnical report at the moment of signing. Each seal generates an **.asice** container verifiable offline that flows into the golden thread (BSA 2022) and the BIM information model (ISO 19650), and pairs with the Declaration of Performance and the Digital Product Passport of Regulation (EU) 2024/3110. It is the same C2PA provenance principle of chapters 13 and 15, applied here not to a synthetic image but to the physical history of a building.

Now the limit, in all its rawness. "Garbage in, garbage on-chain": the seal proves that *that* photo, *that* test result or *that* drawing existed at an instant, was created by a given device or author and has not been altered since. **It does not prove the building is safe**, or that the concrete has the real strength, or that the photographed rebar is the one that stayed in place, or that the test was done honestly. A corrupt inspector can photograph compliant rebar and then pour bad concrete; a laboratory can seal a test that was run but deliberately underloaded; an owner can seal the true drawing of an illegal floor. And it cannot retroactively certify the millions of structures already built and unsealed — Turkey's 13 million

noncompliant buildings. The seal narrows the fraud because it makes later manipulation and fabrication **provable**; it does not guarantee that the physical reality underneath is sound.

The promise and the peril

The promise is a **tamper-proof** works record: that no operator can inspect itself in the shadows, that no laboratory can change the computer's date and that, after a disaster, investigators — like NIST at Surfside — can prove **which records are authentic and contemporaneous** and which were fabricated. The proof of existence does not replace accredited laboratories, independent third-party inspection, the site supervisor's signature or forensic investigation; it makes them **accountable**.

The peril calls for the usual nuances. The first, **false confidence**: "it is anchored, therefore it is safe" is exactly the lie the seal does not authorize — anchoring a datum does not make it true. The second, **concentration and cost**: if sealing every milestone becomes expensive or technically heavy, it can push out honest small builders instead of protecting them. And the third, **capture fraud**: if the corrupted inspector is the one holding the device, they can seal a rigged reality with a seal's full evidentiary force — that is why the forensic GPS and the proof of life bind the author to the place, but do not make them honest. The cryptographic layer makes the cheat provable; it does not make it impossible.

It's time to act

For **materials laboratories and accredited inspection bodies**: seal every test result and every pre-pour inspection at the instant they are produced, so the date and the value can never be touched again. For **developers, builders and engineers**: seal the as-built drawings and the geotechnical report at signing and feed them into the golden thread. For **regulators and building control** (the Building Safety Regulator, CE-marking authorities under Regulation 2024/3110): require the sealed record as part of the golden thread and the Digital Product Passport, and recognize it as evidence. For **collapse investigators** (NIST and its counterparts): treat the dated .asice container as the chain of custody that separates the authentic document from the one fabricated afterward. And for **construction product manufacturers**: anchor each batch's Declaration of Performance, so a later substitution fails the fingerprint check.

As we saw in Adıyaman, the death of dozens of people in a hotel built with river sand and legalized with a purchasable paper shows that giving every construction milestone a sealed proof of how it was built and actually verifying it — never forgetting that the seal does not make safe what is not — is fundamental on this front where concrete hides its own history, time2act.

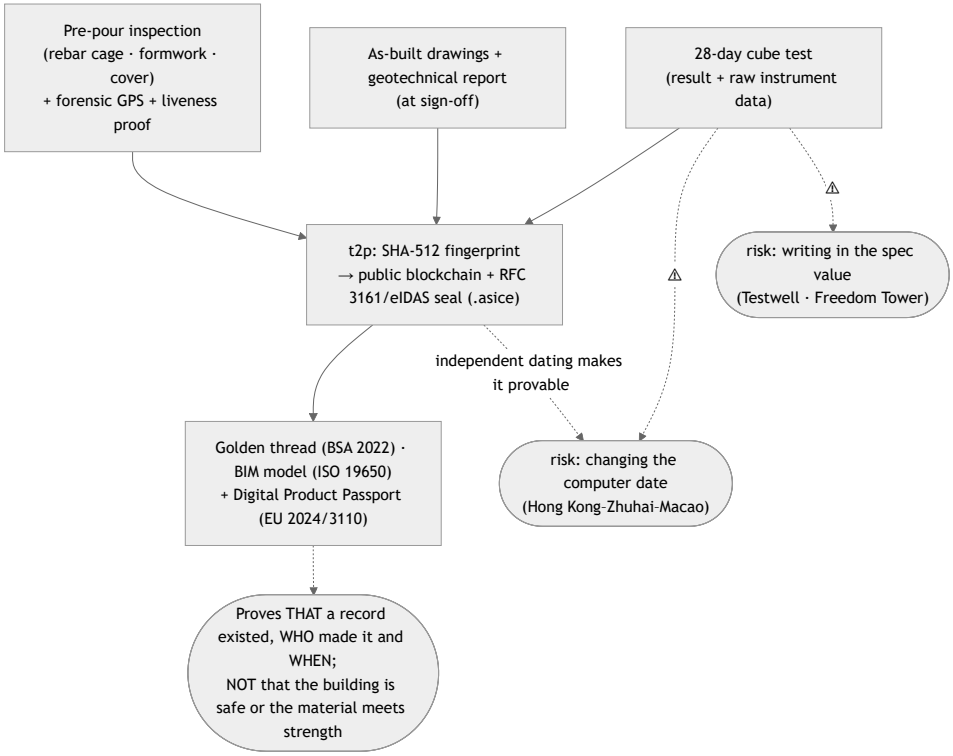


Figure 94.1 — Three irreversible milestones of a build: the pre-pour inspection (rebar, formwork, cover), the 28-day cube test, and the as-built drawings + geotechnical report at signing. At each milestone, t2p computes the SHA-512 fingerprint, adds forensic GPS and proof of life for field captures, seals with RFC 3161/eIDAS and anchors to the public blockchain (.asice); the container flows into the golden thread (BSA 2022) and the BIM model (ISO 19650) and pairs with the Digital Product Passport (Regulation EU 2024/3110). Risk points that independent dating makes provable: changing the computer's date (Hong Kong–Zhuhai–Macao) and writing the specification value (Testwell). Limit: the seal proves that a record existed, who made it and when, not that the building is safe or the material compliant. Sources: 2023 Turkey earthquakes (Isias Grand Hotel); NIST (Champlain Towers South); ICAC (HZMB bridge); Testwell; the Morandi bridge; Building Safety Act 2022; Regulation (EU) 2024/3110.

---

[1] The Dial — *Faulty Construction* (the Turkey earthquake trials; the Isias Grand Hotel; 18-year sentences in December 2024; 75,000 amnestied buildings; some 13 million noncompliant buildings). https://www.thedial.world/articles/news/issue-25/turkey-earthquake-trials ↩

[2] PBS Peril & Promise — *How unenforced building codes worsened the Turkey-Syria earthquake* (July 2023; amnesties, unenforced codes and concrete quality failures). https://www.pbs.org/wnet/peril-and-promise/2023/07/unenforced-building-codes-worsened-the-impacts-of-the-earthquake-that-hit-turkey-and-syria/

[3] NIST — *Champlain Towers South Collapse* (the official federal investigation; 98 dead; initiation by punching shear failure of the pool deck slab; concrete weaker than the code). https://www.nist.gov/disaster-failure-studies/champlain-towers-south-collapse ↩

[4] ICAC Hong Kong — *Consultant director jailed over false HZMB concrete cube test records* (19 workers charged; 384 falsified reports; date manipulation and cube substitution; Albert Leung, 2.5 years in June 2024). https://www.icac.org.hk/en/p/press/index_id_1890.html ↩

[5] The Real Deal — *Concrete tester exec convicted of falsifying safety reports for Yankee Stadium, Freedom Tower* (February 2010; Testwell/Kancharla; the Freedom Tower certified at 12,000 PSI but around 9,000 PSI; 7 to 21 years). https://therealdeal.com/new-york/2010/02/17/concrete-tester-exec-convicted-of-falsifying-safety-reports-for-yankee-stadium-freedom-tower/ ↩

[6] Wikipedia — *Ponte Morandi collapse* (Genoa, August 14, 2018; 43 dead; 59 defendants; "clearly false" inspection documents; self-surveillance via Spea; the €30 million settlement in 2023). https://en.wikipedia.org/wiki/Ponte_Morandi_collapse ↩

[7] Seton UK — *Understanding the Golden Thread: Building Safety Act 2022* (a traceable, version-controlled digital record for higher-risk buildings, post-Grenfell). https://www.seton.co.uk/legislationwatch/article/understanding-the-golden-thread-the-building-safety-act-2022/ ↩

[8] EUR-Lex — *Regulation (EU) 2024/3110* (the revised Construction Products Regulation; in force January 7, 2025; the Digital Product Passport mandatory for CE marking from January 8, 2026). https://eur-lex.europa.eu/eli/reg/2024/3110/oj/eng ↩

[9] NIST — *Report to Congress provides update on Champlain Towers South investigation* (February 2025; experimental work complete, modeling underway, technical presentation in June 2025). https://www.nist.gov/news-events/news/2025/02/nist-report-congress-provides-update-champlain-towers-south-investigation

[10] BSI — *ISO 19650 Building Information Modelling standards* (information management with BIM and the as-built record / digital twin across the life cycle). https://www.bsigroup.com/en-GB/products-and-services/standards/iso-19650-building-information-modelling-bim/

A Part's Pedigree: When a Forged Certificate Can Kill

"Rule books are paper — they will not cushion a sudden meeting of stone and metal."

— Ernest K. Gann, Fate Is the Hunter (1961)

Between January 2019 and July 2023, **Jose Alejandro Zamora Yrala**, aged 38, ran from a home office in Surrey a British aircraft-parts trading company called **AOG Technics**. He sold more than **60,000 jet-engine parts** — worth some £6.9 million — accompanied by **forged release certificates** (EASA Form 1 and FAA Form 8130-3) that fraudulently guaranteed the parts were new and airworthy. Most were for the **CFM56**, the world's best-selling commercial engine, and the falsified pedigree reached airlines such as American Airlines and Ethiopian Airlines. The fraud unraveled in mid-2023, when an airline asked **Safran** to verify an AOG certificate, Safran flagged it as fake and, on August 4, 2023, the UK CAA, the FAA and EASA issued safety alerts. CFM International, the GE Aerospace–Safran joint venture, sued in the London High Court and ended up identifying **126 suspect engines**; the UK Serious Fraud Office secured Zamora's guilty plea in December 2025 and a sentence of **4 years and 8 months** on February 23, 2026, for a fraud valued at £39.3 million [1].

This chapter is not about the silicon of counterfeit chips — that is chapter 70 — nor the generic traceability of goods of chapter 29. Here the object is narrower and **safety-critical**: the **documentary pedigree of a single physical part**. The airworthiness certificate of a turbine blade, the logbook entry, a car's odometer. The question is uncomfortable: if the part is metal that does not speak, and everything that makes it legitimate is a paper that can be forged, how do we make that paper unalterable and attributable? Because here a false document is not a commercial nuisance: **it can kill**.

A paper fraud on flying metal

The AOG case reveals the attack's exact anatomy. Zamora was not making bad parts in a workshop; he manipulated and invented the documents that accompany them. CFM International's lawsuit identified 78 falsified documents covering 52 CFM56 part numbers and two CF6 records, turbine blades included, which reached 68 engines mounted on older Boeing 737s and Airbus A320s; as we have seen, the count of suspect engines rose to 126 [2]. What made the fraud lethal was not a visible metallurgical failure, but blind trust in a certificate: an Authorised Release Certificate with a stamp and a signature no one could check against the part's real origin.

AOG is no isolated case, but the tip of a systemic problem. An FAA estimate cited by the industry calculates that some 520,000 counterfeit or unapproved components reach aircraft every year — close to 2% of the 26 million parts in service — and puts the global market in counterfeit parts above $2 billion a year [3]. Each of these parts travels with a paper pedigree, and each pedigree is only as strong as the difficulty of forging it; today that difficulty is low: a PDF with a logo and a scanned signature.

The same hole, on four wheels: the odometer

The documentary pedigree problem is not exclusive to aviation; it is the same trust hole at scale, affecting millions of consumers. As we have seen, in December 2025 CARFAX estimated that some 2.45 million vehicles with allegedly rolled-back odometers are on U.S. roads, a 14% increase over 2024 and an average loss of some $3,300 per affected car [4]. Digitization has not solved it: rolling back a digital dashboard with a cheap tool is cleaner and faster than dismantling the vehicle.

In Europe the magnitude is comparable. A 2018 European Parliament study found odometer tampering in 5-12% of domestic used-car sales and in 30-50% of cross-border ones, with an estimated economic damage of between €5,600 and €9,600 million a year in the EU [5]. The story is the same as in aviation: a part — the vehicle — that does not speak, and a history — the kilometers, the service book — that can be rewritten. The inspection book with a workshop stamp is, in essence, the humble sibling of the EASA Form 1.

The system the industry is already building

The institutional response exists and moves in the right direction. After AOG, the Aviation Supply Chain Integrity Coalition — Airbus, Boeing, GE Aerospace, Safran, American, Delta, United and StandardAero — was created in February 2024, and on October 9, 2024 it published 13 recommendations centered on three axes: supplier accreditation, digitization of the documents and signatures (FAA Form 8130, EASA Form 1) and better parts traceability [6]. The key word is *digitize*: the coalition recognizes that hand-signed paper is precisely what AOG knew how to exploit.

The regulatory framework has been in place for decades. The FAA's Suspected Unapproved Parts (SUP) Program, established in 1995, requires traceability back to manufacture under 14 CFR Part 21 or an airworthiness determination under Part 43, verified with FAA Form 8130-3 tags, and suspect parts are reported with FAA Form 8120-11 [7]. In Europe, the EASA Form 1 is the equivalent, and EASA maintains a suspect-parts database and issues safety bulletins — like SIB 2023-05 on AOG Technics, of August 2023. On the road, the Federal Odometer Act (49 U.S.C. ch. 327) is enforced by NHTSA in the United States, and in Europe Directive 2014/45/EU and the Belgian Car-Pass system point toward an EU-scale mileage registry agreed in December 2025; as we have seen, the law demands traceability, but making it unalterable is still missing.

Why digitization alone does not save

It must be said without detours: digitizing the EASA Form 1 or the service book **does not by itself prevent the AOG fraud**. A signed PDF is a PDF: it can be touched up, backdated, copied and reissued with an invented signer, exactly as Zamora did. If digitization only replaces the ink stamp with a scanned signature image, we have swapped a forgeable paper for a forgeable file. The same goes for the car's digital dashboard: an electronic kilometer record that can be overwritten is no more reliable than the mechanical needle it replaces.

The underlying problem is not the medium, but the absence of an **independent anchor** fixing *who* declared *what* and *when*, outside the control of whoever might want to rewrite it. A centralized database helps, but it is only as strong as trust in its administrator and as vulnerable as its key. What is needed is for every pedigree event — every certificate, every logbook entry, every odometer reading — to be bound to a certain date and an author in a way that makes any later manipulation **detectable**.

The proof of existence applied

This is where **t2p** adds the missing layer, and it does so on the **pedigree event**, not on the metal. When maintenance creates or transfers a record — an OEM/PMA shop releasing a part, an MRO signing an EASA Form 1 or an FAA 8130-3, a teardown report, a logbook entry or a back-to-birth traceability package — the technician or quality officer captures the document on the device. The **SHA-512 fingerprint** is computed locally, an **RFC 3161/eIDAS timestamp** is applied, the fingerprint is anchored to the public blockchain and an **.asice** container verifiable offline is generated. Adding **forensic GPS** and a **proof of life** binds the human signer to the act. The result is an **unalterable chain of custody**: any later party — an airline's goods-in, the CAMO, a regulator, an auditor or a used-car buyer consulting the sealed book and the odometer photo — can independently verify that *this* certificate or *this* photo existed, unaltered, at *this* moment, signed by *this* author in *this* place. The system sits beneath the 8130-3/Form 1 and the SUP and Car-Pass databases, dovetailing directly with the coalition's recommendation to "digitize documents and signatures + improve traceability." It works alongside the provenance layers of chapters 13 (C2PA) and 15.

Now the limit, in all its rawness. "Garbage in, garbage on-chain": the seal proves **the origin, the time, the author and the integrity** of the certificate, the photo or the logbook entry; **it does not prove that the part is genuine, new, airworthy or properly installed**, nor that the odometer reading or the documented service record is true. A forger can still **seal a lie**: had AOG sealed its own fabricated ARC, the seal would faithfully show that AOG created *this* document at *this* instant — but it would not say that the bolt is really new from Safran and not a salvaged or counterfeit one. A sealed mileage photo indicates that the dashboard showed *this* number at capture, not that it had never been rolled back before. The seal performs no physical inspection, no metallurgical test, no back-to-birth audit, nor does it validate that the issuing organization really held the approval it claims. It is an integrity and provenance layer that makes retroactive manipulation **detectable** and responsibility **traceable**; it never replaces accredited inspection, OEM authentication, regulatory oversight or human engineering judgment.

The promise and the peril

The promise is to close the exact hole AOG exploited: to make **forging a certificate or inventing a signer** leave a detectable trace, and to turn a part's pedigree — or a car's history — into something anyone can verify without trusting the seller's word. In a chain where a turbine blade passes from hand to hand without any unalterable receipt, anchoring the origin protects the honest buyer and the passenger alike.

The peril calls for the usual nuances. The first: **false confidence** — "it is sealed, therefore the part is good" — when the seal only guarantees the document, not the metal; an impeccable .asice on top of a counterfeit bolt is more dangerous than no paper at all, because it inspires trust. The second: **concentration**, if only the big OEMs and MROs can seal and the small shops are pushed out of the market for being unable to. The third: dependence on **honest capture at the origin** — if the first person filling in the Form 1 lies, the seal perpetuates the lie with a certain date. The proof of existence is a pedigree governance layer, not a certificate of the part's goodness.

It's time to act

For the OEMs and the engine industry (GE Aerospace, Safran, CFM, Airbus, Boeing): seal the 8130-3 form at issuance, so that any later touched-up copy fails the fingerprint check. For MROs and CAMOs: verify the sealed fingerprint at goods-in before installation and demand it as a purchasing requirement. For regulators (FAA, EASA, CAA): recognize the sealed pedigree as evidence within the SUP Program and treat the certain date as the difference between a certificate and an after-the-fact alibi.

For the supply chain integrity coalition: make the anchored seal the recommended technical standard for "digitizing documents and signatures." And for automotive regulators and mileage registries (NHTSA, Car-Pass, the future EU registry): seal every odometer reading at inspection, so that any rollback leaves a detectable gap.

As we have seen, Zamora Yrala sold sixty thousand parts with false papers from a home office; only when a Safran officer checked the certificate was the fraud detected. Giving every critical part a pedigree no forger can rewrite without a trace — never forgetting that the seal does not inspect the metal — is, on this front where a false document can kill, time2act.

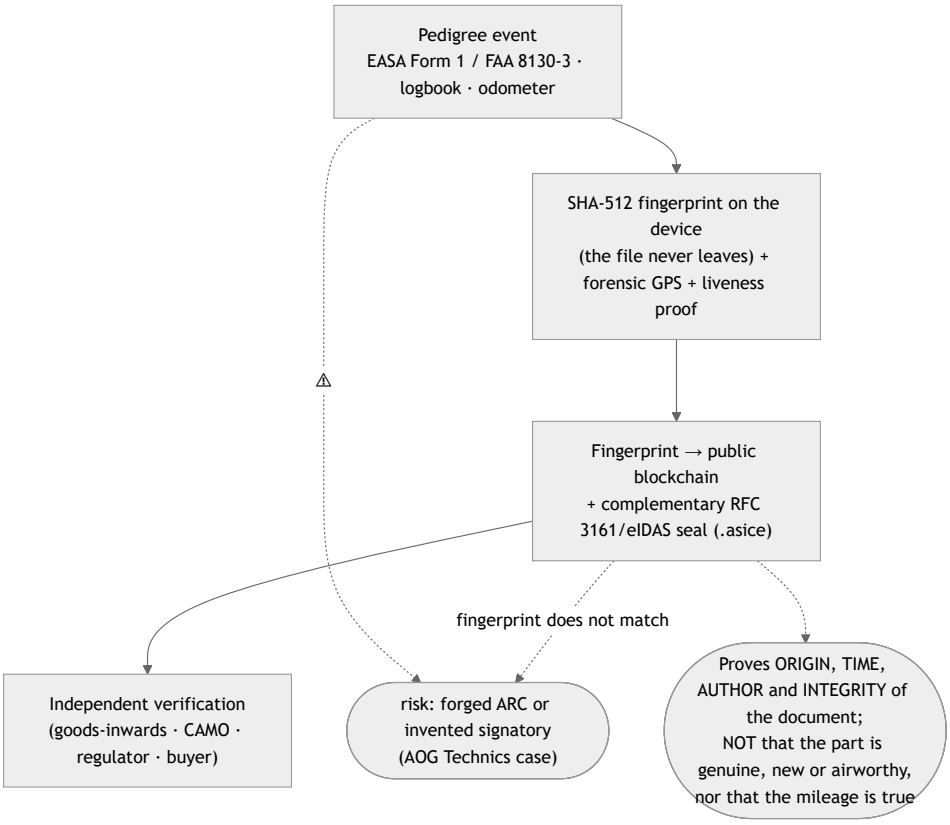


Figure 95.1 — A maintenance event creates a record (EASA Form 1 / FAA 8130-3, logbook entry, odometer reading). The technician computes its SHA-512 fingerprint on the device (the file does not leave), adds forensic GPS and proof of life for field captures, and the whole's fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). Any later party — goods-in, CAMO, regulator, buyer — independently verifies who declared it and when. Limit: the seal proves the document's origin, time, author and integrity; it does not prove the part is genuine, new or airworthy, or the mileage true — garbage in, garbage on-chain. Sources: the AOG Technics case (SFO); CFM International's lawsuit; FAA estimates of unapproved parts; CARFAX and the European Parliament on odometers; the Aviation Supply Chain Integrity Coalition.

---

[1] SFO (GOV.UK) — *SFO secures 4-year prison sentence for aircraft parts fraud* (Zamora Yrala; 60,000 parts; a £39.3 million fraud; forged ARCs; a 4-year-8-month sentence; discovery by Safran). https://www.gov.uk/government/news/sfo-secures-4-year-prison-sentence-for-aircraft-parts-fraud ↩

[2] SFO (GOV.UK) — *SFO case page: AOG Technics Ltd* (the Serious Fraud Office's official case page). https://www.gov.uk/sfo-cases/aog-technics-ltd ↩

[3] Fortune — *Fake components went into 68 jet engines on Boeing 737 and Airbus A320 planes, lawsuit says* (CFM International's suit: 78 falsified documents, 52 CFM56 part numbers, 68 engines, turbine blades). https://fortune.com/2023/09/08/fake-components-68-jet-engines-boeing-737-airbus-a320-planes-lawsuit/ ↩

[4] FAA — *Suspected Unapproved Parts (SUP) Program* (traceability under 14 CFR Part 21/43, Form 8130-3, reporting with Form 8120-11). https://www.faa.gov/aircraft/safety/programs/sups ↩

[5] EASA — *Suspected Unapproved Parts* (the EU database and definitions; cases of forged EASA Form 1; the SIB framework). https://www.easa.europa.eu/en/domains/aircraft-products/suspected-unapproved-parts ↩

[6] EASA — *Safety Information Bulletin SIB 2023-05* (the bulletin on AOG Technics's suspect parts and the recommended quarantine/replacement actions). https://ad.easa.europa.eu/blob/EASA_SIB_2023_05.pdf/SIB_2023-05_1 ↩

[7] GE Aerospace — *Coalition delivers report to help prevent future unapproved parts entering aviation* (the Aviation Supply Chain Integrity Coalition: 13 recommendations, October 9, 2024; digitize 8130-3/Form 1; traceability; supplier accreditation). https://www.geaerospace.com/news/press-releases/coalition-delivers-report-help-prevent-future-unapproved-parts-entering-aviation ↩

[8] Connector Supplier — *Counterfeit Components Ground Airlines* (~520,000 counterfeit or unapproved parts reach aircraft each year, ~2% of 26 million; a market above $2 billion). https://connectorsupplier.com/counterfeit-components-ground-airlines/

[9] PR Newswire / CARFAX — *Odometer rollbacks see dramatic jump of 14% to 2.45 million vehicles* (December 2025; 2.45 million suspect vehicles; +14% year-on-year; ~$3,300 average loss). https://www.prnewswire.com/news-releases/carfax-odometer-rollbacks-see-dramatic-jump-of-14-to-2-45-million-vehicles-302643855.html

[10] European Parliament (EPRS STU 2018-615637) — *Odometer manipulation in motor vehicles in the EU* (tampering in 5-12% of domestic sales and 30-50% of cross-border ones; damage of €5.6-9.6 billion a year). https://www.europarl.europa.eu/RegData/etudes/STUD/2018/615637/EPRS_STU(2018615637_EN.pdf

Emissions: The Sensor That Lied Before the Paperwork Did

"When you can measure what you are speaking about, and express it in numbers, you know something about it; when you cannot measure it, your knowledge is of a meagre and unsatisfactory kind."

— Lord Kelvin, "Electrical Units of Measurement" (1883)

On September 18, 2015, the U.S. **Environmental Protection Agency** (EPA) issued **Volkswagen** a notice of violation of the *Clean Air Act*. Its TDI diesel cars carried a **"defeat device"**: engine-control software that recognized the laboratory test bench — the still steering wheel, the speed, the duration — and only then fully engaged the nitrogen oxide controls; on the road it dialed them down and emitted NOx at up to **40 times** the legal limit. There were some **11 million vehicles** worldwide (about 590,000 in the U.S., model years 2009-2016), and by mid-2020 the scandal had already cost VW more than **$33.3 billion** in fines, buybacks and settlements, including a $2.8 billion criminal penalty and a $1.45 billion civil one [1][2]. The script repeated: in January 2024, the engine maker **Cummins** accepted a civil penalty of **$1,675 million** — the highest ever imposed under the *Clean Air Act* — for undeclared defeat software on more than **630,000** RAM 2500/3500 diesel pickups of model years 2013-2019 [3]. The common thread is raw: the compliance datum the regulator trusted looked truthful, but it had been **fabricated at the moment of capture**. The paperwork was not faked afterward; the sensor or the test was gamed.

This chapter is not about carbon credits or energy savings certificates — that is the business of chapters 21-23, where the asset is a tradable instrument and the risk is double counting or the phantom reduction. Here the asset is something else: the **emissions datum from the sensor and from compliance**, and the risk is the **manipulated or rigged measurement**. The question is precise: how do we fix that *this* reading came from *this* analyzer, in *this* place and instant, without anyone being able to "bias" it between the sensor and the spreadsheet? The answer touches the heart of the book's thesis: when environmental data is worth money and decides sanctions, its **integrity** stops being a technical detail and becomes the battle.

When the test bench is fooled

The Volkswagen case revealed a form of deception that leaves no trace on paper because it is born before the document. The software detected that the car was in a homologated test and engaged a clean mode that was never used on the road; the laboratory reading was **authentic**: there, the vehicle emitted little — but it was not really driving. Thus the fraud endured for years: every certificate was technically correct and internally coherent, and no document audit would have found it, because the document was not lying; the reality it described was.

The Cummins case, eight years later, shows this was no single-brand anomaly. As we have seen, more than 630,000 trucks carried undeclared software, a record civil penalty and a total settlement above $2 billion, with more than $325 million for remedies and a recall of at least 85% of the affected vehicles within three years [3]. In Europe, the **Court of Justice of the EU** has closed the legal door on the "elegant" trick: in July 2022 it ruled that **"thermal window"** software — which only guarantees NOx control between 15 and 33 °C — is, in principle, an illegal defeat device; an earlier judgment, of December 2020, had already declared unlawful the software that recognizes the test cycle [8]. The message is clear: cheating the test is cheating the law, even if the sheet comes out clean.

When the sensor watching the stack is rigged

The test bench trick has an even more direct industrial sibling: manipulating the continuous emissions monitor watching a stack. **Berkshire Power Company** and **Power Plant Management Services** pleaded guilty to felonies against the *Clean Air Act* for intentionally biasing the plant's **continuous emissions monitoring system (CEMS)** and underreporting NOx between January 2009 and March 2011, as already mentioned. The maneuver was of surgical frankness: they introduced a fraudulent adjustment of about **0.5 ppm**, **removed it before** the independent auditors arrived and **put it back** when they left. BPC paid a $2.75 million criminal fine; by March 2017 the parties owed some $7.25 million in fines and penalties [4]. There was no false document anywhere: the CEMS recorded and transmitted, faithfully and in real time, a fabricated figure. The fraud lived in the **blind window** between two audits.

The same hole exists at planetary scale with methane. The IEA's *Global Methane Tracker 2024* estimates close to **120 Mt** of methane from the fossil fuel sector in 2023; but the figures oil and gas companies report to the **UNFCCC** are 50% lower than the IEA's estimate, and those reported to OGMP 2.0 imply about 5 Mt — **95%** lower [5]. The satellite **MethaneSAT** (launched March 4, 2024, communication lost June 20, 2025) found U.S. oil and gas methane emissions **more than four times** higher than the EPA's estimates, with an aggregate loss rate of 1.6% across twelve basins — eight times the industry's 0.2% target [6]. The distance between what is declared and what is measured from the air is, itself, the proof of the problem.

What the law already demands

The regulatory response has hardened on both sides of the Atlantic. The U.S. *Clean Air Act* bans defeat devices and makes it a crime to falsify, tamper with or **knowingly render inaccurate** any required monitoring method — the legal hook of both *Dieselgate* and CEMS manipulation. The **EPA's final methane rule** (signed December 2, 2023, published March 8, 2024), within the *New Source Performance Standards* 40 CFR Part 60 subparts OOOOb and OOOOc, creates a **Super-Emitter Program**: a "super-emitter event" is a release of **100 kg of methane per hour** or more, detected by EPA-certified third parties via satellite, aircraft/drone or ground mobile sensors; the operator must start investigating within **5 days** and report to the EPA within **15 days** [4][6]. The rule is expected to avoid **58 million tonnes** of methane between 2024 and 2038 (a ~80% reduction).

In the European Union, **Regulation (EU) 2024/1787** on the reduction of methane emissions in the energy sector imposes source-level **MRV** (measurement, reporting and verification) with **independent third-party verification**, a satellite-based super-emitter rapid alert mechanism, and equivalent obligations for importers from 2027 [10]. And enforcement already bites: on July 11, 2024, the EPA and the Justice Department announced a **$241.5 million** settlement with **Marathon Oil** for *Clean Air Act* violations at nearly 90 facilities on the Fort Berthold reservation (North Dakota) — a $64.5 million civil penalty, the highest ever applied to a stationary source — cutting some 110,000 tonnes of VOCs and methane [7]. With the *Inflation Reduction Act*'s *Waste Emissions Charge*, every declared tonne of methane has a price. When the datum is worth money, the incentive to dress it up grows.

Why the rule alone does not save

It must be said without detours, because self-deception here would be expensive. A rule that requires an accredited CEMS, a verified MRV and timely reporting does not, by itself, prevent the number that reaches it from being a lie from the origin. As we have seen, the *Clean Air Act* banned defeat devices long before 2015; Volkswagen installed them anyway, and the homologation system, built on trusting the laboratory test, failed to detect them for years. The Berkshire case shows the hole even more nakedly: the entire CEMS regime worked on paper — the sensor transmitted, the regulator received — while the operator slipped the bias in and out around the visits. What fails is not the legal threshold or the duty to report; it is that between the sensor and the spreadsheet there is a **chain of custody no one seals**, and a number can be biased, backdated or substituted without a trace. A compliance report is a PDF; nothing prevents touching up the datum before it becomes a PDF.

The proof of existence applied

This is where **t2p** supplies the missing link, on the artifact that matters: the **raw measurement record at the instant of capture**, before it can be "biased," gamed with the test cycle or silently edited. The CEMS, the methane analyzer or its acquisition system — and likewise the sensor payload of a drone or an aircraft — computes on the device itself the **SHA-512 fingerprint** of each reading or block of raw data; the original file never leaves. It then obtains an **RFC 3161/eIDAS timestamp** and **anchors the fingerprint to the public blockchain**, inside an **.asice** container verifiable offline. On field flights, **forensic GPS** is added so the plume reading is provably bound to a place and moment. The same goes for an engine certification test or a super-emitter notification: the technician or the NGO **seals the test trace** — or the satellite's detection frame — at the moment of producing it. When the operator files the regulatory report, the auditor can verify that the submitted figures **match the sealed raw capture** and that nothing changed between the sensor and the sheet — closing exactly the hole Berkshire exploited by slipping the bias in and out around the audits. It replaces nothing in the existing stack — EPA/CARB

certification, the CEMS's accredited QA, the OOOOb/c and super-emitter program, the MRV of Regulation 2024/1787, the satellite cross-checks (MethaneSAT, IEA, Copernicus): it adds a **tamper-proof chain of custody** from the sensor to the filing. It is the complement, on environmental data, of what chapters 13 and 15 propose for C2PA content provenance.

Now the limit, in all its rawness. "Garbage in, garbage on-chain": the seal proves the origin, time, author and integrity of a measurement record — that this reading came from this instrument, in this place and instant, and has not been altered since. **It does not prove the reading is true** at the point of capture. A defeat device that fools the bench (VW, Cummins) produces a genuinely sealed "clean" result that is, at the same time, empty of meaning; a physically biased CEMS will record, faithfully and verifiably, fabricated low numbers. The seal does not detect that a sensor has been rigged, that a probe has been removed, that dilution air has been introduced or that the engine has entered test-recognition mode; nor does it know whether a facility **never sampled** a leaking component. It does not quantify emissions or judge compliance. It is therefore useless against *Dieselgate*'s central fraud unless combined with independent real-world checks — on-road measurement with portable equipment (PEMS), surprise audits, accredited calibration and satellite and aircraft cross-checks — and it never replaces the regulator's judgment, enforcement or human inspection.

The promise and the peril

The promise is to close the gap between the sensor and the filing: an **unalterable, dated chain of custody** that makes it impossible to bias a number in the blind window between two audits, as Berkshire did, and that gives the honest auditor and the regulator proof that the reported numbers are the ones the sensor produced. In a world where the declared tonne of methane has a price, making the datum's origin tamper-proof protects the honest operator and the public alike.

The peril calls for the usual nuances. The first, do not confuse **integrity with truth**: a false, sealed measurement is still false — and now immutable besides; the seal could breed a false reassurance ("it is anchored, therefore it is good") that corresponds to nothing, above all against the fraud that lives in the sensor, not on paper. The second, the **custodian's capture**: if the party sealing is the same one controlling the rigged sensor, the seal only certifies a false number with more solemnity; that is why independent verification and external cross-checking are non-negotiable. And the third, **concentration and cost**: the seal must strengthen the MRV, not push out small operators or independent observers. The proof of existence is a chain-of-custody layer, not a certificate of environmental goodness.

It's time to act

For **industrial operators and engine makers** (Volkswagen, Cummins, power plants, oil and gas): seal the raw reading of the CEMS and the analyzer, and the trace of every certification test, at the moment of capture, so that no one — not even themselves — can bias it later.

For **regulators** (the EPA, CARB, the European authorities of Regulation 2024/1787): require the filing with a sealed capture and always cross-check it against real measurement — PEMS on the road, satellite on the stack — because no seal can replace inspection.

For the **third-party verifiers and NGOs** of the Super-Emitter Program: seal the detection frame with forensic GPS at the moment of the flight, so that the 100 kg/h notification, as noted, is unassailable.

For **satellite operators and the scientific community** (MethaneSAT, the IEA, Copernicus): maintain the external cross-check no seal can provide, because the datum's integrity and its **plausibility** are two different checks. And for **legislators**: recognize the sealed capture as a valid chain of custody, and remember that it closes Berkshire's hole, not Volkswagen's.

On September 18, 2015, a regulator discovered it had spent years trusting a test the car knew how to recognize. Giving every emissions reading a sealed proof of its origin and its instant — never forgetting that the seal does not make true a measurement rigged at the sensor — is, on this front where environmental data is worth money and decides sanctions, **time2act**.

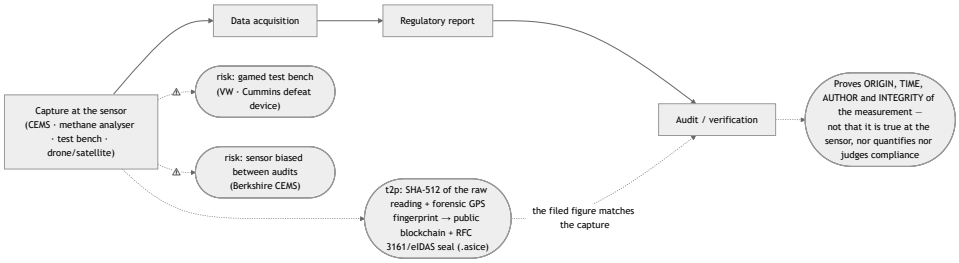


Figure 96.1 — The chain of an emissions datum in four links: capture at the sensor (CEMS, methane analyzer, test bench, drone/satellite flight) → acquisition → regulatory report → audit. t2p computes on the device the SHA-512 fingerprint of the raw reading, adds an RFC 3161/eIDAS timestamp and, in the field, forensic GPS, and anchors it to the public blockchain (.asice); the auditor checks the submitted figure against the sealed capture. Two risk points: the rigged test bench (the VW/Cummins defeat device) and the sensor biased around the audits (Berkshire's CEMS), which the seal does NOT detect. Limit: the seal proves the measurement's origin, time, author and integrity; it does not prove the measurement true at the point of capture, quantify emissions or judge compliance. Sources: EPA (VW, Cummins, Berkshire, Marathon, the methane rule); IEA; MethaneSAT; CJEU (the thermal window).

---

[1] EPA — *Volkswagen Clean Air Act Civil Settlement* (2016-2017 settlements; a $1.45 billion civil penalty; the mitigation fund; ~590,000 U.S. vehicles; the defeat device description). https://www.epa.gov/enforcement/volkswagen-clean-air-act-civil-settlement ↩

[2] Wikipedia — *Volkswagen emissions scandal* (NOx up to 40 times the limit; ~11 million cars; ~$33 billion total cost; the global chronology). https://en.wikipedia.org/wiki/Volkswagen_emissions_scandal ↩

[3] EPA — *United States and California Announce Diesel Engine Manufacturer Cummins Inc. Agrees to Pay Record $1.675 Billion Civil Penalty* (January 10, 2024; the highest civil penalty under the Clean Air Act; 630,000+ RAM trucks; $325 million in remedies; the recall conditions). https://www.epa.gov/newsreleases/united-states-and-california-announce-diesel-engine-manufacturer-cummins-inc-agrees ↩

[4] EPA — *Summary of Criminal Prosecutions: Berkshire Power* (intentional CEMS biasing, 2009-2011; the fraudulent adjustment removed before the audits; a $2.75 million criminal fine for BPC; ~$7.25 million total by March 2017). https://cfpub.epa.gov/compliance/criminal_prosecution/index.cfm?action=3&prosecution_summary_id=2981 ↩

[5] IEA — *Global Methane Tracker 2024, Key Findings* (close to 120 Mt of fossil methane in 2023; UNFCCC ~50% lower; OGMP 2.0 ~95% lower; satellite-detected leak events +50% year-on-year). https://www.iea.org/reports/global-methane-tracker-2024/key-findings ↩

[6] MethaneSAT (EDF) — *New Data Show US Oil and Gas Methane Emissions Over Four Times Higher Than EPA Estimates* (measured emissions >4 times the EPA's estimates; a 1.6% loss rate, eight times the 0.2% target; 30+ flights covering ~70% of U.S. onshore production). https://www.methanesat.org/project-updates/new-data-show-us-oil-and-gas-methane-emissions-over-four-times-higher-epa-estimates ↩

[7] EPA — *EPA and Justice Department Announce $241.5M Settlement with Marathon Oil* (July 11, 2024; a $64.5 million civil penalty, the highest for a stationary source; ~90 facilities; ~110,000 tonnes of VOCs/methane cut). https://www.epa.gov/newsreleases/epa-and-justice-department-announce-2415m-settlement-marathon-oil-reduce-climate-and ↩

[8] CJEU — *Press release on the July 2022 judgment that thermal-window defeat devices are in principle illegal* (PDF; the legal definition of the defeat device; the earlier December 2020 judgment on test-cycle recognition). https://curia.europa.eu/site/upload/docs/application/pdf/2022-07/cp220124en.pdf ↩

[9] EPA — *Key Things to Know About EPA's Final Methane Rule* (fact sheet PDF; 58 million tonnes of methane avoided 2024-2038; a ~80% reduction; the Super-Emitter Program). https://www.epa.gov/system/files/documents/2023-12/key-things-to-know-about-epas-final-rule-for-oil-and-natural-gas-operations.fact-sheet.pdf

[10] European Commission — *New EU Methane Regulation to reduce emissions from fossil fuels* (Regulation (EU) 2024/1787; source-level MRV with third-party verification; satellite super-emitter rapid alert; MRV equivalence for importers from 2027). https://energy.ec.europa.eu/news/new-eu-methane-regulation-reduce-harmful-emissions-fossil-fuels-europe-and-abroad-2024-05-27_en ↩

Environmental Citizen Science: The Sensors from Below, with a Certain Date

"The obligation to endure gives us the right to know."

— Jean Rostand, quoted by Rachel Carson in Silent Spring (1962)

Doñana, one of Europe's most valuable wetlands, was shrinking in silence while around it the strawberry greenhouses sucked the aquifer dry with wells, ponds and kilometers of pipes that no official registry counted. The complaints to Brussels went back to 2009; on June 24, 2021, the Court of Justice of the European Union condemned Spain for failing its obligations: the 2015-2021 Hydrological Plan did not count the illegal extractions and had no measures to prevent the deterioration of the protected habitats [1]. The sentence that sums it up hurts: what the administration does not measure, legally does not exist — even when a national park is being drunk dry.

Chapter 96 dealt with the industrial sensor watched from above: the stack, the official meter, the fraud of the lying sensor. This chapter looks in the opposite direction: the sensors from below, the citizens' own — the air quality device on the balcony, the photo of the dumping, the noise diary, the volunteer's eye on the aquifer. The orbital witness is in chapter 104; the integrity of scientific data, in chapter 63. The key question: how can a citizen's datum weigh enough to move an inspection, a case file or a judgment? The answer: not by pretending to be what it is not — an official report — but by being what no one else can be: a dense, local, persistent chronology, with a certain date and unassailable integrity thanks to **t2p** — the SHA-512 fingerprint computed at the origin, the RFC 3161/eIDAS timestamp and the anchor to the public blockchain, inside an .asice container [5].

What is not measured does not exist

Doñana's paradox is worth unpacking, because it is the pattern of the whole chapter. The European judgment does not punish a pollution event: it punishes a data void. The state did not count the illegal extractions — the infrastructure of wells and pipes draining the groundwater to irrigate strawberries in greenhouses raised on systematic deforestation — and, not counting them, it could not derive protection measures either [1]. Organizations and citizens had been reporting it since 2009; it took twelve years for European justice to say what the territory already knew [1].

The lesson is structural: environmental protection is, before anything else, a question of data. And when the administration does not measure — for lack of means, of will, or through capture by interests — the void can only be filled by organized citizens: neighborhood platforms, land-defense organizations, scientific volunteering. But this filling always hits the same wall: "and how do we know your data is real?" The photo without reliable metadata, the volunteer's spreadsheet, the home sensor's series — all of it is technically impeachable, and the polluter has lawyers to impeach it.

The answer is not to give up on citizen data: it is to armor what can be armored — when it was taken, where, and that no one has touched it since [5]. Metrological quality is defended with calibrations and method; integrity and date, with cryptography. Separating these two planes — and never confusing them — is this chapter's whole secret.

The lagoon that became a person

The Mar Menor changed the whole frame. After decades of degradation, a popular legislative initiative gathered 639,826 signatures, and the result was Law 19/2022, of September 30: the recognition of legal personhood for the lagoon and its basin — the first ecosystem in Europe to become a subject of rights, with the right to exist and evolve naturally, to protection, conservation and restoration [2]. The Constitutional Court upheld its constitutionality, and Royal Decree 90/2025 regulated its governance: a committee of representatives of thirteen members, seven of them from civil society — initially, from the very group that promoted the initiative [2].

The meaning for this book goes beyond the legal milestone: the citizens did not just watch over the ecosystem — they became its legal voice. And a legal person, to defend its rights, needs the same as any other: evidence. Whoever speaks for the lagoon before a court will have to prove the state of the water, the

specific episode, the dumping, the chronology of deterioration or recovery. Citizen science moves from activism to procedural testimony — and a witness needs infallible memory.

This is where the proof of existence closes the circle: every documented sample, every photograph, every series of volunteer measurements, with the SHA-512 fingerprint generated at the origin, the timestamp and the anchor to the public blockchain [5]. Legal personhood grants the standing; the proof of existence, the credibility. One without the other is a right without muscle, or muscle without a right.

The network of low-cost sensors

The global muscle of citizen science already exists. Sensor.Community — born as luftdaten.info — is a network of air quality sensors assembled by citizens, with real-time data and a presence in more than 100 countries [3]; in the Spanish-speaking sphere, networks like AireCiudadano offer particle monitors with an open map and API, and projects like CanAirIO measure PM2.5 with open hardware and a mobile phone [3].

The strength of these networks is the one no official network can have: density and proximity. A sensor on every balcony, next to the school, in front of the factory — a mesh so fine it captures the local variations that the big reference stations, by necessity, average out. And a near-zero marginal cost, which turns maintenance into a collective act.

The weakness is twofold, and it must be said plainly. First, metrological quality: a low-cost sensor is not a reference station — it drifts, gets dirty, is affected by humidity. Second, the chain of custody: who guarantees that the series has not been touched up, that the date is what it claims, that the sensor was where it says?

The right division of labor attacks each weakness with its own weapon. The first is fought with method: cross-calibrations with official stations, documented protocols, hardware transparency [3]. The second, with proof of existence: every batch of readings anchored at the origin — fingerprint, certain date — so that after-the-fact touch-ups become impossible and the series' gaps are dated too [5]. A humble sensor with an armored series is worth more, before a technician or a judge, than an excellent sensor with an editable spreadsheet: the certainty of who, what and when weighs as much as precision.

Noise: the proof in two voices

The most everyday use case is also the most misunderstood. When excessive noise is reported, the city council must make an inspection visit, and the measurement is done by municipal technicians or the local police with a homologated, calibrated sound level meter, drawing up a report with a presumption of veracity and full evidentiary value — date, time, the offender's identity, the type of noise and the measurement [4]. The report of an acoustic expert with homologated equipment is powerful evidence in litigation [4]; the indicative indoor limits hover around 35-45 decibels by day and 25-30 by night [4]; and the Ombudsman has had to remind administrations to provide the means to check complaints [4].

The system's hole is structural: noise is intermittent and the inspection is punctual. The bar that blasts music into the small hours goes quiet the night the technician comes; the machine that vibrates every dawn rests on the day of the visit. The neighbor who suffers it holds the information the report cannot capture — the recurrence — but their phone recording is not an official report, and never will be.

The division of labor, again: the citizen does not replace the official sound level meter; they provide the chronology that gets it there on time. A sealed noise diary — dated entries night after night, with fingerprint and attested integrity [5] — does not prove decibels, but it proves persistence: dates, times, durations, a pattern no punctual inspection can deny and that lets the official visit be scheduled when it counts. The official measurement takes the photo; the sealed diary makes the film. And a case file needs both.

The proof of existence applied

A balcony, a thirty-euro sensor and a data series no lawyer can erase.

1. The series no one can touch up

The citizen air quality sensor anchors every batch of readings at the origin: SHA-512 fingerprint, timestamp, public blockchain [5]. When the neighborhood platform presents three years of data to the city council or the court, the argument about "doctored data" ends before it starts: the series is exactly what it was, dated gaps included. Legitimate criticism is confined to where it belongs — calibration, method [3] — which is exactly the argument citizen science can win with transparency. It is the counterweight to chapter 96: there, the stack's sensor watched by the regulator; here, the citizens watching the watchman.

2. The dumping photographed

The slurry in the stream, the black smoke at dawn, the rubble on the vacant lot: the environmental complaint lives or dies by the quality of its trail. A series of photographs sealed in the moment — fingerprint, certain date, location — builds the progression: not a loose image that "could be from any day," but a chronology with attested integrity that the environmental authority can verify before opening a case [5]. And the same shield protects the honestly accused: against the old photo recycled in a smear campaign, the absence of a contemporaneous seal speaks for itself.

3. The noise diary

The neighbor who suffers the machine every dawn records, night after night, the episode: date, time, duration, an indicative recording — all sealed on their device [5]. After three months, they present the city council with a chronology of attested integrity that turns their word into a verifiable pattern, and the inspection with the homologated sound level meter — the only one that produces a report with evidentiary value [4] — is scheduled for when the pattern says it will happen, not when the bar expects it. The sealed diary does not measure decibels: it makes the official measurement arrive at the moment of truth.

4. The lagoon's voice

The guardians of the Mar Menor — the civil society sitting on the ecosystem-person's committee of representatives [2] — can give their monitoring cryptographic memory: documented water samples, photographed episodes, community sensor series [3][5]. When the lagoon's right "to exist and evolve naturally" [2] has to be defended in court, its legal voice will carry verifiable chronologies, not militant memories. And the same logic holds for Doñana: the citizen chronology of extractions and dry-outs, sealed year after year, is the permanent counterweight to that hydrological plan that "did not count" [1].

The limits, honestly: the seal attests to when and what was recorded, not that the sensor measured well — calibration and method remain the scientific battle [3]; citizen data is not an official report or an expert opinion, and must not pretend to be [4]; and environmental monitoring has an ethical frontier: facts, facilities and episodes are documented — never the neighbors' lives. Without method, the seal armors noise; without ethics, it armors a siege.

The promise and the peril

The promise is a citizenry that no longer depends on the administration's willingness to measure [1]: armored sensor series [5], ecosystems with a legal voice and a memory [2], official inspections that arrive on time [4] and case files that hold. Democratic environmental monitoring: measure, certify, act.

The perils demand the same clarity.

The first is **the neighborhood-police**. The line between documenting pollution and surveilling people is thin. The camera points at the stack, not at the neighbor's kitchen; at the stream, not at whoever strolls beside it. Environmental facts are documented; faces, license plates and other people's lives have no place, and the fingerprints-only design legitimizes no dissemination [5].

The second is **false equivalence**. Sealed does not mean accurate. A poorly calibrated datum with a certain date is still a bad datum — more dangerous, in fact, because it looks solid. If citizen science sells the seal as a certificate of quality, it will lose the credibility it seeks: the seal attests to integrity and moment; accuracy is attested by calibration and method [3].

The third is **instrumentalization**. An industrial or political interest can also disguise itself as a "neighborhood platform" and flood the file with biased sealed chronologies. The antidote is not cryptographic but methodological: transparency about who measures, with what, how and under what protocol. The seal guarantees that no one has touched up the data — not that it was collected in good faith.

And the fourth is **the administrative excuse**. If citizen data serves for the administration to wash its hands — "the neighbors are already measuring" — we will have lost. The CJEU condemned Spain because the state did not count [1], and the Ombudsman demands means to inspect [4]: citizen science is a counterweight and complement to the public duty to measure, never its substitute.

It's time to act

Nature cannot testify; those who watch over it can — if they carry dates that are proof against doubt.

For citizens and environmental organizations: set up sensors with documented protocols [3], seal the series and the photographs at the origin [5], and think in long chronologies — persistence is your strength.

For city councils: open channels to receive sealed citizen chronologies and schedule the official inspection — the report with the homologated sound level meter [4] — when the pattern indicates; the Ombudsman is already reminding you [4].

For water authorities and environmental agencies: count everything — Doñana's lesson is that the data void is paid for in Luxembourg [1] — and treat the anchored citizen series as what it is: qualified evidence that deserves checking.

For citizen science platforms: build in anchoring by default [3][5]: it is the difference between a pretty map and a case file that holds.

For the guardians of ecosystems with rights: give the Mar Menor's voice cryptographic memory [2] — every sample and every episode, with a certain date.

For legislators: recognize the sealed citizen chronology as sufficient grounds to trigger inspections, and fund the public measurement network — the counterweight should not have to play the lead.

Doñana took twelve years to get a judgment because the extractions "were not counted" [1]; the Mar Menor needed 639,826 signatures to have a voice [2]. Let the next lagoon not have to wait so long: give it witnesses with a certain date. For the land that cannot testify, **time2act**.

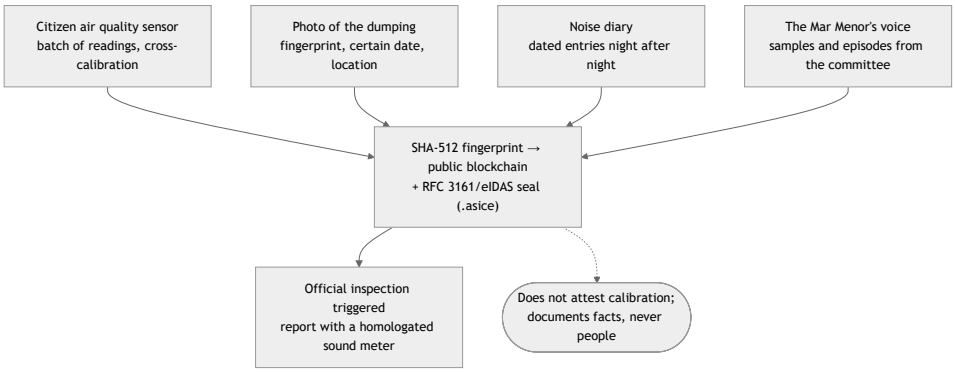


Figure 97.1 — Citizen science with proof of existence: the balcony sensor, the photo of the dumping and the noise diary are anchored at the origin (SHA-512 fingerprint, RFC 3161/eIDAS timestamp, public blockchain, .asice) and become chronologies of attested integrity that trigger the official inspection (a report with a homologated sound level meter, with a presumption of veracity), guide the expert and give memory to the legal voice of ecosystems (Mar Menor, Law 19/2022). Limits: the seal does not attest calibration or replace the official report; facts are documented, not people. Sources: CJEU (Doñana, 2021); BOE (Law 19/2022); Sensor.Community; the Ombudsman.

---

[1] El Ágora Diario — *Europa condena a España por no proteger el agua y hábitats de Doñana* (the CJEU judgment of June 24, 2021; illegal extractions not counted in the 2015-2021 Hydrological Plan; complaints since 2009); Greenpeace — *Doñana, agua sin ley* (wells, ponds and pipes; strawberry greenhouses; deforestation). https://www.elagoradiario.com/agua/tribunal-justicia-europa-condena-espana-agua-donana/ · https://es.greenpeace.org/es/en-profundidad/sos-acuiferos/donana-agua-sin-ley/ ↩

[2] BOE — *Law 19/2022, of September 30, recognizing legal personhood for the Mar Menor lagoon and its basin* (the first ecosystem in Europe; rights to exist, evolve, protection, conservation and restoration; a popular initiative with 639,826 signatures); Araoz & Rueda — *El TC reconoce la personalidad jurídica del Mar Menor*; BOE — *Royal Decree 90/2025* (committee of representatives: 13 members, 7 from civil society); Newtral — *Qué significa que el Mar Menor haya adquirido personalidad jurídica*. https://www.boe.es/buscar/doc.php?id=BOE-A-2022-16019 · https://www.araozyrueda.com/el-

tribunal-constitucional-reconoce-mar-menor-sujeto-de-derechos/ · https://www.boe.es/diario_boe/txt.php?id=BOE-A-2025-2577 · https://www.newtral.es/mar-menor-personalidad-juridica/20220923/ ↩

[3] Sensor.Community / World Air Quality Index — *Real-time citizen air quality sensor network (formerly luftdaten.info)* (more than 100 countries; low-cost citizen sensors); AireCiudadano — an open network of particle monitors with a map and API; SciStarter/CanAirIO — open PM2.5 sensors. https://aqicn.org/network/luftdaten/es/ · https://aireciudadano.com/ · https://scistarter.org/crea-tu-propio-sensor-de-calidad-del-aire-con-cana ↩

[4] Perito Judicial — *Ruidos molestos y su reclamación: la clave está en la medición* (municipal inspection with a homologated, calibrated sound level meter; reports with a presumption of veracity and full evidentiary value; indicative indoor limits 35-45 dB day / 25-30 dB night); COITTA — *La importancia de la prueba pericial acústica*; the Ombudsman — *Comprobaciones ante las denuncias de ruidos molestos* (providing the means to check complaints). https://peritojudicial.com/ruidos-molestos/ · https://www.aagit.org/la-importancia-de-la-prueba-pericial-acustica-y-la-proteccion-juridica-frente-al-ruido-como-agente-contaminante/ · https://www.defensordelpueblo.es/resoluciones/habilitar-los-medios-para-poder-realizar-comprobaciones-necesarias-cuando-se-denuncien-ruidos-molestos/ ↩

[5] t2p — product manual (SHA-512 fingerprint at the origin — sensor or device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; real-time batch series; dated gaps; content on the device and voluntary disclosure). See chapters 13 and 15; the industrial sensor, in chapter 96; the orbital witness, in chapter 104. ↩

The Product Passport: Who Swears to the Recycled Content, and When

"If names be not correct, language is not in accordance with the truth of things."

— Confucius, Analects, XIII, 3

On June 4, 2024, **Volvo Cars** unveiled on its new electric SUV, the **EX90**, what the brand and its partner **Circulor** called "the world's first battery passport," almost three years before the European mandate makes it obligatory. Built over five years with the British blockchain company Circulor, at a cost of roughly **$10 per vehicle** for fifteen years of monitoring, the passport lets the owner scan a **QR code** on the driver's door pillar — or open the Volvo app — and see the battery's chemistry, the origin and production journey of its cobalt, nickel, lithium ions and graphite, the **recycled content percentage** and the embedded **carbon footprint** of the whole pack. Circulor claims to have already collected more than **20 billion traceability points** across **150 million cells**, working with suppliers representing roughly **52% of world cell production** [5][6]. The case is vivid precisely because it exposes the underlying tension of the digital product passport: the document faithfully records and dates everything the supply chain declares, but the EX90's own proof-of-concept lineage — the 2023 Global Battery Alliance pilot with Audi and Tesla — openly admitted that "the quality of the data presented has not been fully verified and was partially complemented with reasonable estimates" [7]. A beautifully structured passport can, then, carry an inflated or unverified recycled-content claim. It is exactly the greenwashing risk the EU's own 2020 sweep quantified: close to **42% of green claims** by European companies were false or misleading.

This chapter deals with the **life cycle passport of a specific physical object** across its whole life, distinct from chapter 29 — the supply chain, which is the custody and provenance of inputs as goods pass from one actor to another. Here the subject is not the flow of materials, but **a single battery**: its recycled content percentage, its reparability and disassembly data, its **state of health** in use and its **second-life** and end-of-life status — claims that persist and update *after* the product has been sold. The question is not whether a beautiful passport can be built — Volvo has already built one — but how to make every figure it contains **provable** and not merely declared.

The faithful document of a datum no one has checked

A digital passport is, in essence, a container of declarations: structured, machine-readable and accessible with a QR. Its virtue — gathering everything in a single verifiable place — is also its Achilles' heel, because the document is exactly as good as the data that goes into it. The **Global Battery Alliance** pilot, presented in Davos on January 18, 2023 with Audi, BASF, CATL, Glencore, Tesla and Volkswagen, said it with rare frankness: the data "has not been fully verified" and was "partially complemented with reasonable estimates" [7]. That sentence, buried in a triumphant press release, is the chapter's key. The passport does not manufacture truth; it carries what it is given.

And what it is given are **falsifiable figures with legal consequences**. From August 18, 2031, the European battery regulation imposes minimum recycled content thresholds — **cobalt 16%**, **lead 85%**, **lithium 6%**, **nickel 6%** — which in a second step from 2036 rise to **cobalt 26%**, **lithium 12%**, **nickel 15%** [2]. These are precise percentages the passport must declare, and each one is a temptation: one percentage point of "extra" recycled cobalt cheapens compliance and improves the image. Virgin cobalt relabeled as recycled at the smelter leaves no visible trace in the document; the passport, docile, records it as recycled. This is the difference between a system that **structures** the information and one that **proves** it.

The state of health that survives the sale

There is a second family of declarations that makes the product passport radically different from simple supply chain traceability: those that **update while the object lives**. Circulor announces fifteen years of **state of health** monitoring for the EX90's battery [6]. This figure — the capacity the pack retains relative to when new — will determine the price of a secondhand EX90 in 2032, whether a battery is fit for a **second life** as stationary storage or must go to recycling. It is a claim that lives and changes after the sale, and its economic value is immense.

Precisely for that reason it is manipulable. The state of health comes from the battery management system (the BMS); a diagnostic export is a file, and a file can be touched up. An unscrupulous seller has every incentive to dress up the capacity of a worn battery before reselling the car; a second-life operator, to declare fit a cell that is not. The passport, again, will obediently record whatever figure it is handed. The promise of a transparent secondhand and second-life market — one of the great environmental gains of the whole circular economy — depends on these readings being **authentic**, not merely **present** in the document.

The regulatory framework: the passport as law, not favor

Europe has turned the passport into an obligation, not a voluntary gesture. The **battery Regulation (EU) 2023/1542**, repealing the 2006 directive and applying since February 18, 2024, makes the **battery passport mandatory from February 18, 2027** for every electric vehicle battery, every industrial battery over 2 kWh and every light means of transport (LMT) battery placed on the European market, accessible via a QR linked to a unique identifier; batteries placed on the market after that date without a registered passport **cannot legally be sold**, and customs authorities get direct access to the registry at the border [1][3]. The **carbon footprint declaration** already began for electric vehicle batteries on February 18, 2025 and extends to industrial ones over 2 kWh on February 18, 2026; it follows the ISO 14040/44 method and the European PEF and **must be verified by a notified body — it cannot be self-certified** [3].

Above it, the **Ecodesign for Sustainable Products Regulation (EU) 2024/1781** (ESPR), in force since July 18, 2024, elevates the digital product passport to a **horizontal tool** that will extend to textiles, iron and steel, aluminum, furniture and tires under the 2025-2030 work plan, as we have seen, adopted on April 16, 2025; the battery is its legally binding pilot [8]. The **CEN-CENELEC JTC 24** committee is developing the horizontal standards (unique identifiers, data carriers, the physical-digital link, interoperability, access rights), and the Commission-funded **CIRPASS / CIRPASS-2** projects define the cross-sector data model. One detail reveals the legislator's caution: the supply chain **due diligence** obligations were postponed two years, from August 18, 2025 to August 18, 2027, through Regulation (EU) 2025/1561, the "stop-the-clock" omnibus [4]. The law demands precise figures; the system that must make them reliable is still being built.

Why the passport alone does not save

It must be said bluntly. A digital product passport, however well structured and full of QRs, **is only as reliable as the will of whoever pours the data into it**. A recycled content certificate posted in the passport is a PDF; nothing prevents the percentage from having been inflated before entering it. A state of health export is a file; nothing in the document proves it reflects the real pack and not a touched-up BMS reading. A carbon footprint declaration can follow the PEF method to the letter and at the same time feed on false activity data. And the notified body's verification, though mandatory for the carbon footprint, does not extend to every state of health reading or every recycling certificate across fifteen years of life.

The passport solves the problem of **having** the information in an accessible, standardized place. It does not solve the problem of **attesting** that the information was true and has not been touched. What is missing is a layer that fixes, for each underlying evidentiary artifact — the recycler's certificate, the notified body's report, the BMS export — **who issued it, when, and that it has not been altered since**. Without this layer, the GBA pilot was right: passport data is often "estimated" and "not fully verified," and the document blesses it all the same.

The proof of existence applied

This is where **t2p** supplies the missing piece, and it does so on the **evidentiary act** behind each claim, not on the database entry. At the moment a supplier or a manufacturer generates a primary artifact — a recycler's batch certificate with the recycled cobalt or lithium-ion content, a notified body's carbon footprint verification report, a BMS state-of-health diagnostic export at the end of the first life, a disassembly or second-life inspection report, or a photo of a batch of material at a mine or a sorting plant — the device locally computes the **SHA-512 fingerprint** (the file **never leaves** the device), obtains an **RFC 3161 timestamp conforming to eIDAS**, anchors it to the public blockchain and packages it in an

.asice container verifiable offline. For field capture — a mine, a recycling line, a dismantling bay — the **forensic GPS** and a **proof of life** bind the record to a real place and a real person at a real moment.

The passport's QR can carry, next to the human-readable claim, a **verifiable link** to that sealed artifact. Anyone — a customs officer at the border, the buyer of a used EX90 reading the second-life state of health, an auditor checking the recycled content threshold — can confirm **offline** that *this* recycling certificate or *this* state-of-health reading existed, unaltered, on *this* date and from *this* author, before the claim was published. It complements the regulated stack — the Article 7 notified body verification, the GBA's data assurance rulebook, the JTC's data carrier standards, Circulor-style blockchain traceability — by giving each underlying document a tamper-proof provenance trail, without ever replacing verification. It is, on the product's life cycle, the sibling of the C2PA mark and SynthID of chapters 13 and 15: those seal **what a model produces**; t2p seals **what an object has lived**.

Now the limit, in all its rawness. **"Garbage in, garbage on-chain"**: the seal proves the origin, the instant, the author and the integrity of a life cycle artifact; **it does not prove that the claim it contains is true**. It does not confirm that the recycled cobalt was really recycled and not virgin metal relabeled at the smelter, or that a sealed state-of-health export reflects the real pack and not a manipulated BMS, or that a life cycle analysis used honest activity data, or that a "second-life" battery is actually safe to redeploy. A falsified recycled content certificate, once sealed, becomes a **tamper-proof** falsified certificate. t2p fixes provenance and unalterability; it does not replace the notified body's conformity assessment, the due diligence audit, the physical testing of the material or the human judgment that decides whether the sealed content is legitimate.

The promise and the peril

The promise is a circular market where a product's lineage is verifiable and not merely published: being able to prove that this recycled percentage, this carbon footprint and this state of health were declared by this actor on that date, and that no one has touched them since. For the buyer of a used EX90, for the auditor — as we have seen — and for the customs officer, making the origin unalterable turns a marketing promise into a claim with consequences.

The peril calls for the usual nuances. The first, do not confuse integrity with truth: a passport full of well-sealed artifacts with inflated data is still a deceptive passport, now with a patina of technological rigor that can lull the auditor — "it is anchored, so it must be true" — when the seal says nothing about the content's truth. The second, concentration: if sealing and passporting pays off only for the big manufacturers — Volvo's $10 per vehicle scales poorly for a small repair shop — the standard can push small repair and second-life actors out instead of strengthening them. And the third, privacy and trade secrets: the passport exposes chain data that can be sensitive, and transparency must be balanced with the legitimate right not to reveal everything. The proof of existence is a life cycle governance layer, not a certificate of environmental goodness.

It's time to act

For **manufacturers and brands** (Volvo, its automotive peers and, soon, those of textiles and electronics): seal every evidentiary artifact at the moment of generating it — recycling certificate, carbon footprint report, state-of-health reading — and link it from the passport's QR, so every figure has a certain date and an author. For **recyclers and notified bodies**: seal the batch certificates and the verification reports at issuance, so the auditor can confirm them offline without depending on the registry's good faith. For **regulators and customs authorities** (the Commission under the ESPR, JTC 24, customs with direct access from 2027): recognize the sealed artifact as evidence of the provenance the passport already requires, and treat the certain date as the difference between a declaration and an alibi fabricated afterward. For **passport platform providers** (Circulor and its competitors): anchor the fingerprint of each underlying document, not just structure its content. And for **buyers and second-life operators**: demand the verifiable link to the sealed state of health before paying for a used battery, so the reading is authentic and not merely present.

In Gothenburg, on June 4, 2024, a car left the factory with the world's first battery passport and a footnote admitting the data was not "fully verified." Giving every product a sealed proof of every claim it makes across its life — never forgetting that the seal does not make true the recycling someone declares — is, on this front where an object must answer for what it says about itself, time2act.

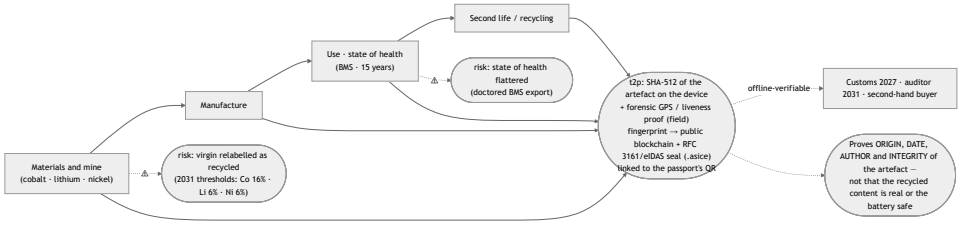


Figure 98.1 — A battery's life cycle in four links: materials and mine → manufacture → use and state of health → second life / recycling. Each link generates an evidentiary artifact (recycling certificate, the notified body's carbon footprint report, the BMS export, the disassembly report); t2p computes its SHA-512 fingerprint on the device and anchors it to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), with forensic GPS and proof of life for field capture. The passport's QR links the claim to the sealed artifact, verifiable offline by customs, auditors and secondhand buyers. Limit: the seal proves the artifact's origin, date, author and integrity; it does not prove the recycled content real, the state of health untouched or the battery safe. Sources: Volvo EX90 / Circulor; the Global Battery Alliance; Regulation (EU) 2023/1542; the ESPR (EU) 2024/1781.

---

[1] EUR-Lex — *Sustainability rules for batteries and waste batteries (Regulation (EU) 2023/1542 summary)* (the passport from 2027 via QR and unique identifier; recycled content targets from August 18, 2031: Co 16%, Pb 85%, Li 6%, Ni 6%; customs access to the registry). https://eur-lex.europa.eu/EN/legal-content/summary/sustainability-rules-for-batteries-and-waste-batteries.html ↩

[2] DigiProdPass — *EU Battery Regulation: Recycled Content Targets (2031-2036)* (the 2031 percentages — Co 16/Pb 85/Li 6/Ni 6 — and the 2036 ones — Co 26/Pb 85/Li 12/Ni 15 — the falsifiable claims the passport carries). https://digiprodpass.com/blogs/eu-battery-regulation-recycled-content-targets ↩

[3] Circularise — *EU battery passport regulation requirements* (the phased calendar: EV carbon footprint Feb 18, 2025 / industrial Feb 18, 2026; the passport Feb 18, 2027; notified body verification; EV/industrial >2 kWh/LMT scope). https://www.circularise.com/blogs/eu-battery-passport-regulation-requirements ↩

[4] Minespider — *EU Battery Due Diligence Obligations postponed* (due diligence postponed from Aug 18, 2025 to Aug 18, 2027 via the "stop-the-clock" Regulation (EU) 2025/1561; guidance extended until Jul 26, 2026). https://www.minespider.com/blog/eu-battery-due-diligence-postponed-what-it-means-for-companies ↩

[5] Electrek — *Volvo EX90 launches world's first EV battery passport* (Jun 4, 2024; five years of development with Circulor; ~$10/vehicle for 15 years; the QR on the door pillar; traces materials, recycled content and carbon footprint). https://electrek.co/2024/06/04/volvo-ex90-launch-worlds-first-ev-battery-passport/ ↩

[6] IEEE Spectrum — *Volvo's Battery Passport Traces the Supply Chain* (more than 20 billion data points across 150 million cells; ~52% of the world's cell makers; 15 years of state-of-health monitoring). https://spectrum.ieee.org/volvo-ex90-battery-passport ↩

[7] Global Battery Alliance — *World's First Battery Passport Proof of Concept* (Davos, Jan 18, 2023; Audi/BASF/CATL/Glencore/Tesla/VW; the data "has not been fully verified and was partially complemented with reasonable estimates"). https://www.globalbattery.org/press-releases/global-battery-alliance-launches-world%E2%80%99s-first-battery-passport-proof-of-concept/ ↩

[8] European Commission (Green Forum) — *Implementing the Ecodesign for Sustainable Products Regulation (ESPR)* (the ESPR in force Jul 18, 2024; the digital product passport as a horizontal tool; the 2025-2030 work plan adopted Apr 16, 2025). https://green-forum.ec.europa.eu/implementing-ecodesign-sustainable-products-regulation_en ↩

[9] Hogan Lovells — *Digital Product Passports in the EU: expansion under ESPR and lessons from the Battery Passport pilot* (legal analysis distinguishing the battery pilot from the ESPR's horizontal rollout; the delegated acts calendar). https://www.hoganlovells.com/en/publications/digital-product-passports-in-the-eu-comprehensive-expansion

[10] CIRPASS — *Standardisation gaps and roadmap (v1.2, July 2024)* (the DPP data model, unique identifiers, data carriers and the JTC 24 standardization gaps — the architecture the seal plugs into). https://cirpassproject.eu/wp-content/uploads/2024/07/CIRPASS-standardisation-gaps-and-roadmap-V1.2.pdf

Waste: Where the Bin Really Goes

"Nobody wonders where, each day, they carry their load of refuse."

— Italo Calvino, Invisible Cities (1972), "Leonia"

Every night, in millions of kitchens, the same choreography repeats: rinse the yogurt cup, fold the cardboard box, and take it all down — the yellow bin for packaging, the blue for paper, the brown for organic waste. It is a gesture of pure trust: none of us has ever seen what happens next. We trust that the gesture means something, that the journey ends well, that "it gets recycled" means it gets recycled.

This chapter's question is exactly that: where does the bin really go? And the urgency is not rhetorical [1]: Europe exports waste on an industrial scale — Turkey alone received 39% of European exports, 12.4 million tonnes in one year — and the trajectory of these shipments has been dark enough that the Union has had to intervene: Regulation (EU) 2024/1157 on waste shipments, generally applicable from May 21, 2026, subjects exports of all plastic waste to prior written notification and consent, bans exporting waste for recovery to countries outside the OECD and imposes electronic processing of shipments [1].

Between your bin and the final "recycled" there is a chain — collection, transfer, licensed manager, treatment, often a truck or a ship — made of changes of hands and one-sided papers. When all goes well, no one looks at it; when it does not, we find out late: the "recycling" that was dumping, the certificate signed without treatment, the green label with nothing underneath.

This chapter deals with the greatest collective act of faith of our time. The new product's passport — its recycled content, its second life — is in chapter 98; rigged emissions, in chapter 96; the origin of goods, in chapter 29; the chain of custody of relief — the same logic, inverted — in chapter 50; citizen evidence, in chapter 97. The key question: how does a system that moves millions of gestures and billions in fees stand on delivery notes no one can verify? The answer: by giving every change of hands a certain date with **t2p** — the fingerprint-time-anchor chain of chapter 13 [3] — from the weighbridge to the final treatment certificate. The Regulation's new rules [1] are worth whatever the delivery note that documents them is worth; the verifiable trace is what makes them count.

The waste chain — and where it breaks

The chain, as it is: the producer — the factory, the shop, the household — generates the waste; collection takes it to a transfer station, where flows are consolidated; the licensed manager routes it to treatment — recycling, recovery or disposal; and often there is a shipment in between: a truck to another region, a container to another continent. The rules already require identification and tracking documents for each step, and the new European Regulation wants them electronic and with prior consent when they cross borders [1].

The breaking points have familiar names in the sector. Cascading subcontracting: the licensed manager who delegates to another, who delegates to a third — and the papers dilute at every step. The final treatment certificate signed far from the treatment. The mass balance that does not add up: more tonnes enter the plant than leave it treated, and the difference evaporates into incomplete reports. The export declared as "recovery" that ends up in a landfill on the other side of the world. And, crowning it all, greenwashing: the "100% recyclable," the "zero landfill" — claims no consumer can check.

The evidentiary key is this whole book's: all these holes pass between one-sided delivery notes — modifiable, losable, without an opposable certain date. Extended producer responsibility schemes pay by declared tonnes; the authorities inspect samples; and the citizen, who sustains the system with their gesture and their fees, can verify nothing at all.

The proof of existence applied

The same chain, link by link — and every change of hands with a certain place, time and weight [3].

1. The container with place and time

At every change of hands — loading the truck, entry to the transfer station, unloading at the manager's — a photo with t2p's forensic GPS: a genuine fix, accuracy of about 50 meters, maximum age of about 30

seconds; a simulated location blocks certification [3]. The container that leaves a municipality and "turns up" where it should not stops being a suspicion: it is a chain of dated captures with a hole exactly where someone diverted it. And the sealing of the export container, filmed with the date and time overlaid [3], travels as evidence with the ship.

2. The sealed mass balance

The weighbridge records — what enters, what leaves, when — certified in batches with chained fingerprints: each closing, a fingerprint that incorporates the previous one [3]. A plant's mass balance becomes a dated series no one can rewrite: if a thousand tonnes enter one month and six hundred leave treated, the difference has a date and demands an explanation. It is the same logic the state already imposes on invoices: the chain of fingerprints as accounting normality.

3. The final treatment certificate

The document that closes the circle — "this waste was treated thus, here, on this day" — scanned and sealed with a certain date, linked by fingerprint to the delivery notes and captures of the whole chain [3]. The producer who pays their extended responsibility knows what they paid for; the authority that authorizes a shipment [1] can demand the destination's anchored certificate; and the honest manager — which is most of them — has in their own trace the defense against the unfair competition of those who only pretend to treat.

4. The green claim, frozen

The "made with 70% recycled material," the "zero waste to landfill," the corporate sustainability page: captured with t2p's certified browser — full image and source code, with a date [3]. The day reality does not square with the slogan, that day's slogan exists, exact and dated. And the bridge to chapter 98: the product passport declares recycled content — the waste trace is the only way that declaration has somewhere to come from.

And the limits, for there are limits: the seal attests to custody — place, time, weight, document — not the quality of the treatment: whether the recycling is truly recycling must be said by inspections and analyses. The mass balance can deceive if the first weighing is already false — the chain is worth as much as its starting point, which is why it must start at the entry weighbridge. And the trace belongs to the cargo and the documents, never to the people: proportionality toward plant workers rules, as in all of this book's labor chapters.

The promise and the peril

The promise: extended responsibility fees paid on verified tonnes, not declared ones; international shipments [1] with files the regulator can audit without traveling; honest managers competing on traceability instead of suffering the competition of fraud; and the citizen's gesture — the nightly bin — finally answered with a chain that can be shown.

And the perils, one by one.

The first is **the seal as makeup**. Certifying false delivery notes with a certain date only makes the fraud more elegant. The trace is worth what its first link is worth: calibrated weighbridges, captures in the moment, inspections that check physically. Cryptography documents; it does not replace control.

The second is **the burden on the small manager**. The municipal recycling point and the small plant cannot absorb infinite bureaucracy. The gesture must be the usual one — the ten-second photo at the weighbridge [3] — and the documentary bulk, automatic; if the trace becomes a barrier to entry, only the big players will remain.

The third is **surveillance of the plant**. Documenting loads cannot become surveilling workers. Weighbridges, containers and documents are captured; people, never — the line of all this book's labor chapters.

And the fourth is **the techno-solutionism of the bin**. The best trace in the world does not reduce one gram of waste. The hierarchy is the eternal one: reduce, reuse — and only then recycle well. The proof of existence makes the third step honest; the first two are production and consumption decisions no hash makes for us.

It's time to act

Let the gesture of millions have the trace it deserves.

For citizens: your gesture deserves a trace — demand verifiable data from your municipalities and packaging schemes, not brochure percentages [1].

For managers and plants: the photo at the weighbridge and the batch balance [3] — your honesty, provable, is your best contract.

For producers and extended responsibility schemes: pay — and charge — on tonnes with a certain date; the anchored treatment certificate as a condition of payment.

For environmental authorities and prosecutors: the new Regulation [1] gives you electronic processing; demand it anchored — the verifiable shipment file is the difference between authorizing and blessing.

For legislators: let the Spanish rollout of the Regulation and the announced deposit-and-return scheme for packaging be born with a neutral certain date as standard — it would be the world's first waste system verifiable from the weighbridge to the certificate.

Let each night's yogurt cup have the journey we promised it — and let whoever diverts it leave fingerprints. For the gesture of millions that deserves to be true, **time2act**.

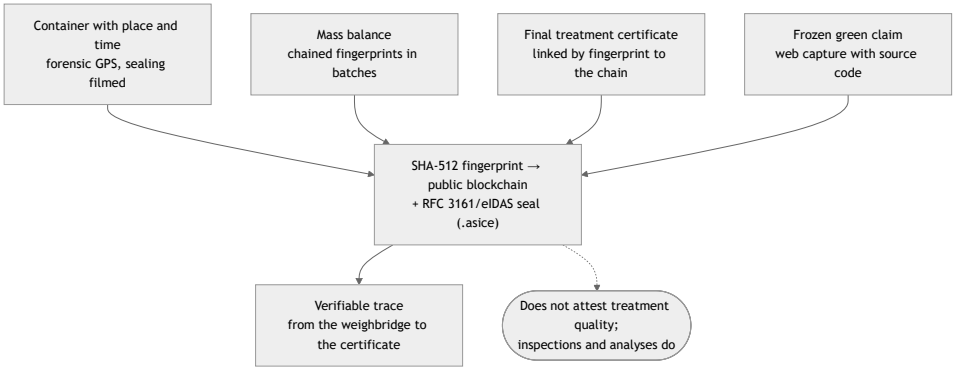


Figure 99.1 — The waste chain with proof of existence (sealed with t2p — chapter 13): the container with place and time at every change of hands (forensic GPS; the sealing filmed with the date overlaid), the mass balance sealed in chained batches (dated inflows vs. outflows), the anchored final treatment certificate and the frozen green claim (web capture with source code). Context: Regulation (EU) 2024/1157, applicable from 21-5-2026 — prior notification for all exported plastic waste, a ban on non-OECD countries; Turkey received 39% of European exports (12.4 Mt). Limits: the trace attests to custody, not treatment quality; reduce and reuse first. Sources: OJEU/BOE; RETEMA.

---

[1] OJEU/BOE — *Regulation (EU) 2024/1157, of April 11, 2024, on shipments of waste* (repeals Regulation 1013/2006; generally applicable from May 21, 2026; exports of all plastic waste to third countries subject to prior written notification and consent; a ban on exporting waste for recovery to non-OECD countries; electronic processing); RETEMA — *Nuevo Reglamento de traslados de residuos*; Residuos Profesional — *Análisis del Reglamento (UE) 2024/1157* (context: Turkey received 39% of European waste exports — 12.4 million tonnes, 2022 data). https://www.boe.es/buscar/doc.php?id=DOUE-L-2024-80606 · https://eur-lex.europa.eu/legal-content/es/ALL/?uri=CELEX:32024R1157 · https://www.retema.es/blog/nuevo-reglamento-de-traslados-de-residuos-la-normativa-que-pone-en-jaque-la-industria-europea · https://www.residuosprofesional.com/reglamento-relativo-traslados-de-residuos/ ↩

[2] The general waste framework (identification and tracking documents for shipments; extended producer responsibility; the waste hierarchy). See chapter 98 for the digital product passport and chapter 96 for emissions.

[3] t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; video with date/time overlay; document scanner up to 50 pages; certify any file in batches — only the fingerprint travels; web page capture with full image + source code; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the product passport, in chapter 98; citizen science, in chapter 97. ↩

The Land and the CAP: The Verifiable Logbook

"The soil is the great connector of lives, the source and destination of all. It is the healer and restorer and resurrector."

— Wendell Berry, The Unsettling of America (1977)

The tractor stopped on the field margin, the phone in a soil-stained hand — and a notification from the administration: the satellite "has detected no agricultural activity" on the declared parcel; a georeferenced photograph must be submitted within a deadline, or the aid is at risk. The farmer looks at the field he sowed six weeks ago — the field that is right there, with the green lines of newly sprouted wheat — and understands what has just happened to him: he now has to prove the obvious to an algorithm.

The scene is becoming routine because the common agricultural policy — the European Union's largest budget — has changed its way of watching [2]: from in-person inspection by sampling, to satellite monitoring. The Sentinels of the Copernicus program scan every declared parcel, and the systems classify whether what was declared — the crop, the pasture, the fallow — can be seen from above; when they cannot conclude — clouds, small parcels, late crops — the burden comes down to earth: a georeferenced photo the farmer takes with his phone [2].

This chapter treats the farm as a subject of proof. The orbital witness and its manipulations are in chapter 104; food fraud — what is really on the plate — in chapter 101; parcel geolocation against deforestation, in chapter 29; citizen sensors, in chapter 97; the farm's waste, in chapter 99. The key question: when the eye that distributes the aid is an orbital algorithm, what does the farmer who really did sow defend himself with? The answer: with his own chronology — the logbook sealed in the moment and in the field, the photo with a neutral certain date — through **t2p**: fingerprint on the device and a certain time anchored to the chain (chapter 13) [3]. The symmetry of chapter 104, seen from below.

The logbook that is coming — and the satellite that is already here

The second piece on the board has an exact calendar [1]: the digital farm logbook — the CUE, the electronic record of each farm's operations: plant protection treatments, fertilization, irrigation and other practices — has been available voluntarily since January 1, 2024 and, after successive postponements, will be mandatory on January 1, 2028 for the largest holdings: more than 30 hectares of crops, more than 5 of irrigated land or more than 0.1 of greenhouse. Before that, an intermediate milestone: the electronic record of plant protection treatments becomes mandatory on January 1, 2027 — paper ends on December 31, 2026 [1].

This book's reading — and it is now the third time it happens: first invoices with VeriFactu, then the working-time record, now the land: the state demands digital records of operations, and watches with orbital eyes. The proof-of-existence architecture is imposing itself on agriculture... but with an asymmetry: the satellite belongs to the administration, the apps belong to the administration, and the logbook will live on platforms the administration consults. The farmer supplies the data; the proof is held by others.

And the technical nuance that changes everything: the satellite gets it wrong too. Persistent clouds, small parcels, late or atypical crops, field margins — the system that "detects no activity" may be looking badly [2]. When this happens, the honest farmer finds himself in the most Kafkaesque position in this book: proving to an algorithm that his field exists.

The proof of existence applied

The logbook and the photo, yes — but sealed in the moment, in the field, and in the farmer's hands [3].

1. The logbook entry, in the moment and in the field

The plant protection treatment, the fertilizing, the irrigation — noted when they are done, where they are done: the logbook entry sealed with t2p with the field's forensic GPS — a genuine fix, accuracy of about 50 meters, maximum age of about 30 seconds; a simulated location blocks certification [3]. The difference from the logbook "reconstructed on Sunday at the table" is evidentiary and it is sanitary: the record made in the moment is the one that counts when the inspection comes, when residue shows up in the product or

when safety intervals are questioned. And when the CUE becomes mandatory [1], the farmer's sealed version will be his cross-check against any platform discrepancy.

2. The photo the satellite demands — with a neutral certain date

The georeferenced photograph the monitoring system requests [2], taken with t2p: certain place and time, with the certificate and its QR code [3]. It is not just completing the formality: it is the farmer's answer carrying the same evidentiary rank as the algorithm's question — evidence verifiable by third parties, not one more image inside an app that does whatever it does.

3. The defense against the eye that errs

The decisive symmetry: the series of dated photos of the crop — the sowing, the sprouting, the growth — made with four gestures across the season [3]. When the satellite "does not see" what is there — the clouds, the small parcel, the late crop [2] — the farmer does not start from zero against the notification: he has an anchored chronology proving the agricultural activity before anyone questioned it. It is chapter 104 seen from below: the orbital witness is powerful, but not infallible — and whoever works the land deserves to be able to contradict it with facts.

4. The aid file

The application, the sealed logbook, the season's photos, the responses to the requests — grouped and exportable as an .asice package verifiable offline [3]: for the appeal against an aid reduction, for the in-person inspection, for the traceability the grain buyer or the cooperative is starting to ask for. And the bridge to the neighboring chapters: what the field produces enters the chain of chapter 101, and from it comes the waste of chapter 99 — the verifiable farm is the first link of both.

The limits, stated clearly: the seal attests to operations, places and dates — it does not replace the official logbook or the administration's platforms: it cross-checks and complements them [1]. It does not make the practices good — the plant protection limits and the aid conditions are what they are [2]. And the countryside's digital divide is real: the tool must be the ten-second gesture with the phone already in the pocket, and the cooperatives and the agronomists, the allies of those who do not take to it.

The promise and the peril

The promise: farmers who answer the satellite's requests with chronologies of their own [2][3]; logbooks that count as evidence because they were written in the moment [1]; aid that is collected and defended with dated facts; and a farm-level traceability the market will begin to pay for — grain with a pedigree is no longer a utopia: it is the natural continuation of chapters 29 and 101.

The perils, faced head-on.

The first is **the surveilled countryside**. Orbital monitoring plus the digital logbook can make the farmer the most audited profession in the country. The answer is this book's: let the trace serve both directions — if the administration watches from above, the farmer must be able to answer from below with the same rank of proof; surveillance without symmetry is submission.

The second is **the bureaucracy that expels**. Every new requirement weighs more on the small farm. The CUE's calendar already recognizes it — mandatory only for the large ones, for now [1]; the tool must subtract work, not add it: the note in the moment is faster than the Sunday reconstruction.

The third is **the algorithm as judge**. If the satellite's notification is treated as truth by default, the agricultural presumption of innocence dies. The monitoring system is a filter, not a verdict [2]; its discrepancies must be resolved with ground evidence — and the burden of proof cannot fall entirely on the farmer.

There remains the fourth, the most tempting: believing this is enough. No seal fixes farm-gate prices, the drought or generational renewal. The proof of existence makes the aid fair and the farms defensible; agricultural policy is made — and must be made well — by people.

It's time to act

Let the field that is there be provable from there.

For farmers and stockbreeders: the note in the moment with forensic GPS [3] — the logbook made in the field is worth twice the one reconstructed; and the season's photo series is your insurance against the eye that errs [2].

For cooperatives and agronomists: you are the layer that makes all this possible for the small farm — integrate sealing into your logbook and application services [1][3].

For the agricultural administration: monitoring saves you inspections [2]; give the symmetry back — accept the farmer's anchored chronology as a qualified cross-check of the satellite's discrepancies, and let your apps issue verifiable evidence, not captive evidence.

For buyers and industry: the farm with its own trace is the first step of your chain — chapters 29 and 101 are its continuation; pay for it what it is worth.

For legislators: the CUE will come [1]; let it come with the proof in the hands of whoever generates it — the logbook sealed by the farmer, with a neutral certain date, as a record with full evidentiary value.

For ten thousand years the peasantry has fed the world without anyone asking it to prove it; now that they ask from space, let it at least have the tools. For the field that is there — and that can be proven — **time2act**.

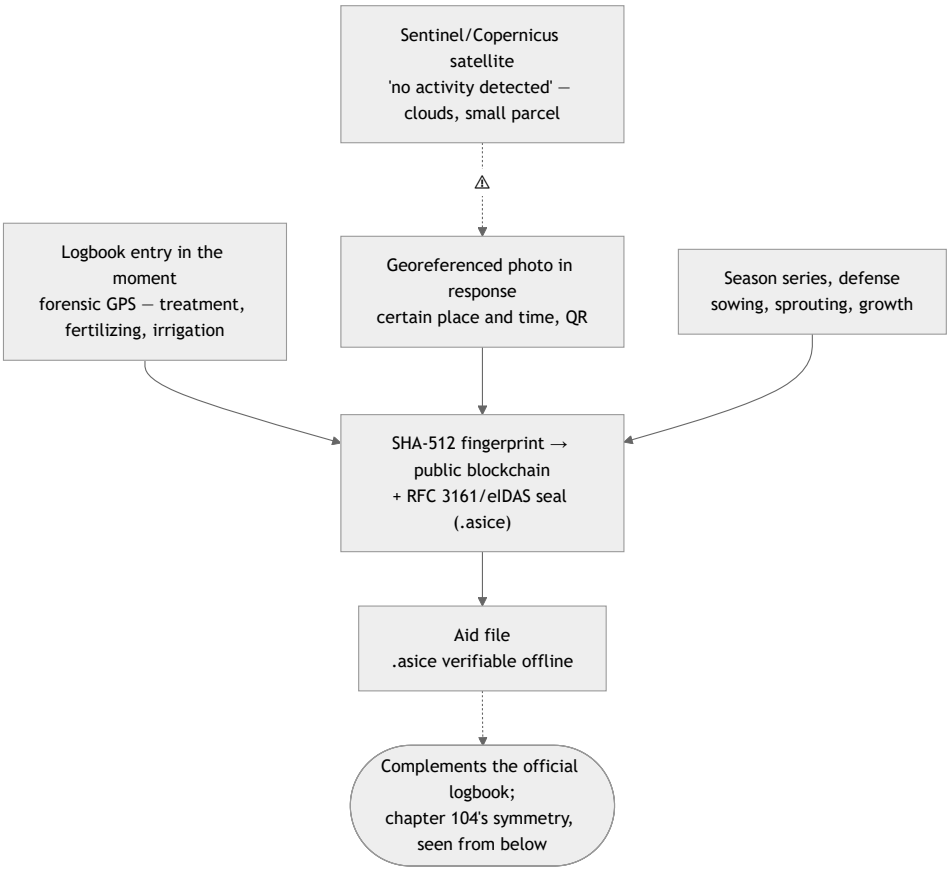


Figure 100.1 — The land with proof of existence (sealed with t2p — chapter 13): the logbook entry sealed in the moment and in the field (forensic GPS), the georeferenced photo the monitoring demands — with a neutral certain date — the season series as a defense against the satellite that errs (chapter 104's symmetry, seen from below) and the aid file in .asice. Context: the CAP watched by Sentinel/Copernicus; the digital logbook (CUE) voluntary since 2024, electronic plant protection records from 1-1-2027, the CUE mandatory on 1-1-2028 (>30 ha). Limits: it complements the official logbook, does not replace it; the countryside's digital divide demands cooperatives and agronomists as allies. Sources: FEGA/CUE rules; CAP 2023-2027.

---

[1] Agrodigital — *El cuaderno digital de explotación será obligatorio a partir del 1 de enero de 2028* (the CUE voluntary since 1-1-2024; postponements; mandatory from 1-1-2028 for holdings >30 ha of crops, >5 ha irrigated or >0.1 ha greenhouse; the electronic plant protection record mandatory from 1-1-2027 — paper until 31-12-2026); FEGA/Agronews — entry-into-force clarifications; Government of Aragon — *Cuaderno digital (CUE)* (content: plant protection operations, fertilization, irrigation and other practices). https://agrodigital.info/2025/01/23/el-cuaderno-digital-de-explotacion-sera-obligatorio-a-partir-del-1-de-enero-de-2028/ · https://www.agronewscastillayleon.com/el-fega-publica-aclara-la-entrada-en-vigor-del-cuaderno-de-explotacion-y-el-plan-de-abonado/ · https://www.aragon.es/pac-2026/cuaderno-digital-cue ·

https://www.fega.gob.es/sites/default/files/files/document/10.2023.03.15-12.45h_blq-04_inmaculada-bravo_cue.el-cue-y-el-riego.pdf ↩

[2] CAP 2023-2027 — the monitoring system (Sentinel/Copernicus imagery over the parcels declared in the SIGPAC; requests for georeferenced photographs from the farmer when the system cannot conclude). The FEGA's and the European Commission's general framework. See chapter 104 for the integrity of the orbital witness. ↩

[3] t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; background certification — the field without coverage; certify any file in batches; document scanner; projects for a season's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the satellite, in chapter 104; food, in chapter 101; waste, in chapter 99. ↩

Food Fraud: When the Label Lies About What You Have Eaten

"Tell me what you eat and I shall tell you what you are."

— Jean Anthelme Brillat-Savarin, Physiologie du goût (1825), aphorism IV

On the night of July 8-9, 2024, Italy's Carabinieri Command for Health Protection (NAS), within the Europol-coordinated OPSON framework, raided eighteen warehouses in the Apulia region and seized some **42 tonnes of adulterated oil** already bottled and labeled as "Italian extra virgin olive oil," plus some 71 tonnes of an "oily substance," worth close to one million dollars. The detail that explains everything: investigators found **623 liters of chlorophyll** being added to cheap oils to fake the green color of authentic extra virgin, with forged "extra virgin" labels and **1,145 excise stamps** suspected of being counterfeit. Seven people were charged with criminal conspiracy, adulteration of food for sale, fraud in public and military supplies, and adulteration for export [3]. It all flowed from Operation Omegabad of November 2023 (Guardia Civil + Carabinieri + Europol), which seized **260,000 liters** of oil — cheap "lampante" cut and sold as extra virgin in Ciudad Real, Sicily and Tuscany — arrested eleven suspects and confiscated €91,000 in cash [4]. The wave tracked the record wholesale price of extra virgin, about €10.50/kg in early 2024, after the Andalusian drought cut harvests almost in half.

This chapter tackles a question distinct from its neighbors'. Chapter 29 covers the generic supply chain — custody, logistics and provenance of goods; chapter 72, illegal extraction and the conflict origin of a raw material. Here we speak specifically of the **authenticity and origin of food**: adulteration (cutting oil with lampante or sunflower, honey with sugar syrup), species substitution (horse for beef) and **false geographical claims** of protected designation. The question is simple and ancient: *is this product what the label says?* And the book's thesis fits here: against a fraud that lives on **false paperwork** — labels, certificates, stamps — the proof of existence fixes who declared what, where and when, with no possibility of backdating.

A business of some $40 billion that lives on the origin premium

Food fraud affects roughly 1% of the world industry and costs between $10,000 and $40,000 million a year [7]. The EU's Agri-Food Fraud Network logged 758 suspicions in 2023, 26% more than the year before, with meat substitution, honey adulteration and mislabeled oil at the top [7]. The annual sales value of EU geographical indication products was €74.8 billion in 2017, close to 7% of food and drink sales, and a protected-designation product sells on average at twice the price of a comparable uncertified one [6]; the more valuable the designation, the more tempting it is to counterfeit.

Operation OPSON XIII (results announced in October 2024) confiscated some 22,000 tonnes of food worth about €91 million, dismantling eleven networks and reporting 278 people; in Italy 42 tonnes of fake extra virgin and 623 liters of chlorophyll were found, and in Spain 45,000 liters of oil and 120,000 cans of tuna [1]. **Operation OPSON XIV** (announced in 2025) raised the figure to some €95 million — 11,566 tonnes of food and 1.4 million liters of drinks — with 631 people reported and thirteen organized crime groups dismantled in 31 countries [2].

The honey that is not honey and the horse that was beef

Oil is not the only front. In the European coordinated action "From the Hives" (samples from November 2021 to February 2022, results published March 21, 2023), **147 of 320 consignments of imported honey — 46%** — turned out suspicious of adulteration with sugar syrups, an enormous jump from the 14% recorded in 2015-17. The calculation is transparent: real honey cost on average about €2.32/kg imported and rice syrup between €0.40 and 0.60/kg; the EU imports some **175,000 tonnes a year**, close to 40% of consumption [5]. So one in two consignments arrived under suspicion of not being real honey.

And no survey is complete without the episode that shook Europe. In the 2013 horse meat scandal, Findus beef lasagna contained more than 60% — up to 100% — horse meat. At the Paris trial of April 16, 2019, the former head of Spanghero, Jacques Poujol, received two years (eighteen months suspended) and a €100,000 fine; the Dutch trader Jan Fasen, two years and a permanent ban from the meat sector, for the

sale of more than **500 tonnes of horse meat** channeled through Comigel to 28 companies in thirteen countries [9]. A label can lie across half of Europe before anyone notices.

The law of origin: a reinforced registry, but made of paper

The regulatory response has been substantial. The EU's geographical indications regime was revised with **Regulation (EU) 2024/1143**, in force since **May 13, 2024**, which merges the PDO, PGI and GI rules for food, wines and spirits into a single registration framework and reinforces the protection of indications used as an ingredient and sold online [8]. **Regulation (EU) 2017/625** on official controls gives the legal basis for controls along the chain and for DG SANTE's **Agri-Food Fraud Network (FFN)**, connected to the RASFF alert system. The oil standards (Regs. 2022/2104 and 2022/2105) define the categories — extra virgin, virgin, lampante — and the honey Directive 2001/110/EC was reinforced in 2024 by Directive (EU) 2024/1438. In the United States, the **FSMA's** preventive controls require protection against economically motivated adulteration (EMA); the canonical example is cheap oil sold as 100% olive oil [10]. And OPSON supplies the operational muscle every year.

The problem is that this whole architecture rests on **documents**: certificates of conformity, batch numbers, laboratory reports, excise stamps and designation labels. A document, as we saw in Apulia, can be forged, cloned, recycled or backdated; the 2024/1143 registry indicates which designation deserves protection, but it does not guarantee that this certificate is genuine or that this analysis covers this batch.

Why the certificate alone does not save

It must be said bluntly. A PDF certificate of conformity, a printed batch number or a signed laboratory report prove nothing if they can be altered without a trace. As we saw in Apulia, the fraudsters created the "extra virgin" labels and the excise stamps to walk the lie past the controls. A genuine certificate for a small batch can cover a much larger shipment diluted afterward; an authentic report on an honest batch can be attached, its date retouched, to an adulterated one; and if the excise stamp is forgeable, the whole chain of guarantees cracks. What is missing is not more paper: what is needed is an unalterable anchor that fixes who issued each authenticity artifact, over which batch and at what instant, so it can be neither cloned nor backdated.

The proof of existence applied

This is where **t2p** supplies the missing layer, on the paperwork and the field act, not on the chemistry. It works in three phases. First, at batch creation, the producer captures geotagged photos and video with proof of life of the olive grove, the hive or the bottling line; the device computes the SHA-512 fingerprint on the spot — the file never leaves — anchors it on the public blockchain, adds a complementary RFC 3161 timestamp conforming to eIDAS, and the forensic GPS binds the act to the claimed PDO/PGI territory, fixing where and when the batch was produced before any blending can occur. Second, the authenticity artifacts are sealed: the accredited laboratory's report (NMR/IRMS profile, isotope ratio, pollen for honey, fatty acids for oil), the certificate of conformity, the batch number, the excise documents and the PDO/PGI label art, each with its fingerprint and sealed in an **.asice** container verifiable offline and bound to a batch identifier. Third, any inspector, importer or court can verify, offline, that the certificate, the report and the origin photos existed then, were issued by that body and have been neither altered nor backdated — deactivating the forged labels, cloned stamps and recycled certificates seen in OPSON and Omegabad. It is the complement, on the label and the document, of what chapters 13 and 15 propose with C2PA for the provenance of images and certificates, and it can feed the digital product passport of chapter 98.

The limit is clear: "garbage in, garbage on-chain." The seal proves the origin, time, author and integrity of a captured artifact; it does not check that the sealed content is true, legitimate or safe. It cannot say whether the oil in the bottle is really extra virgin, whether the honey is cut with rice syrup or whether the geographical claim is chemically true: only laboratory analysis determines that (NMR, IRMS, DNA, pollen). A dishonest producer can still seal an authentic photo of the wrong olive grove, fingerprint a genuine certificate that covers less batch than is shipped, or seal correct paperwork for an adulterated product. It does not stop downstream substitution and is worth nothing if the certifying body itself is corrupt; it is an evidentiary, anti-backdating tool on top of — never in place of — laboratory authentication, official controls, accredited certification and human judgment.

The promise and the peril

The promise is an agri-food ecosystem where origin and documentary guarantee are **verifiable and not backdatable**: being able to prove with a certain date that that certificate, that report and those photos of the olive grove existed, were issued by the legitimate body and have not been touched. In a sector where the designation premium sells at double and fraud costs up to $40 billion a year, **as we have seen**, making the origin unalterable protects the honest producer, the prudent importer and the consumer.

The peril calls for the usual nuances. First, do not confuse **integrity with truth**: false, sealed paperwork is still false — and now immutable besides; the seal could breed a false reassurance ("it is anchored, therefore genuine") with no correspondence to the real chemistry. Second, **dependence on the certifying body**: if the regulatory council or the laboratory is captive or corrupt, the seal eternalizes the lie with an appearance of rigor. And third, the **digital divide**: imposing an advanced seal could expel the small cooperative member or the global South producer without an accredited laboratory or an affordable tool, concentrating the quality market even further.

It's time to act

For **PDO/PGI consortia and regulatory councils**: seal the batch number, the label art and the certificate at issuance, so a cloned label like Apulia's fails verification. For **accredited laboratories and control bodies**: anchor the analytical report (NMR/IRMS, pollen, fatty acids) and the certificate of conformity at the moment of issue, binding them to the batch identifier. For **producers and cooperatives**: capture the geotagged field evidence with proof of life and forensic GPS at the harvest, at the hive or on the bottling line, before any blending. For **inspectors and customs** (NAS, Guardia Civil, OPSON, OLAF): verify offline the authenticity and date of the stamps and certificates, instead of trusting them by eye. And for **regulators and DG SANTE**: recognize the sealed artifact as evidence of the required documentation and feed it into the alert network — as we have seen with Reg. 2024/1143.

In Apulia, one summer night, seven people were filling bottles of cheap oil dyed with **623 liters** of chlorophyll and sticking on fake labels and stamps to pass it off as Italian extra virgin. Giving every honest batch a sealed proof of its origin, its date and its certificate — never forgetting that the seal does not make true what is declared in it — is, on this front where the label can lie about what we eat, time2act.

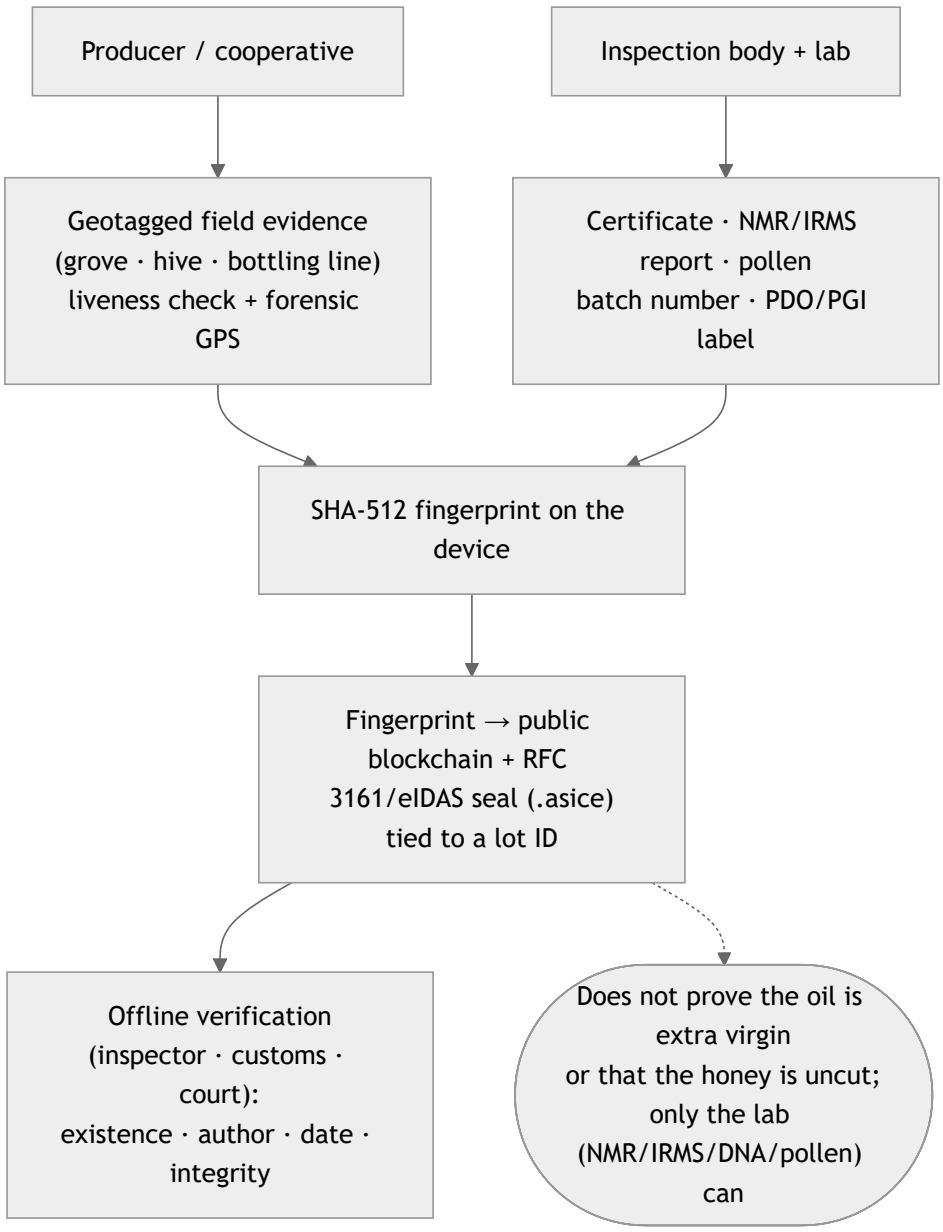


Figure 101.1 — The producer captures the geotagged field evidence with proof of life and forensic GPS (olive grove, hive, bottling line) and computes its SHA-512 fingerprint on the device; in parallel, the control body and the laboratory seal the certificate, the analytical report, the batch number and the PDO/PGI label. The bundle's fingerprint is anchored on the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), bound to a batch identifier. Any inspector, customs office or court verifies it offline: existence, author, date and integrity. Limit: the seal proves the paper and the field act, not that the oil is extra virgin or the honey uncut — the laboratory says that (NMR/IRMS/DNA/pollen). Sources: OPSON XIII and XIV (Europol); the Apulia case; Operation Omegabad; "From the Hives"; Reg. (EU) 2024/1143.

---

[1] Europol — *EUR 91 million worth of counterfeit and substandard food seized in Europe-wide operation* (Operation OPSON XIII; 22,000 t of food, €91 million, 11 networks; Italy: extra virgin and chlorophyll; Spain: oil and tuna). https://www.europol.europa.eu/media-press/newsroom/news/eur-91-million-worth-of-counterfeit-and-substandard-food-seized-in-europe-wide-operation ↩

[2] Europol — *Counterfeit and substandard food worth EUR 95 million seized in global operation* (Operation OPSON XIV; €95 million, 11,566 t, 631 reported, 13 organized crime groups, 31 countries). https://www.europol.europa.eu/media-press/newsroom/news/counterfeit-and-substandard-food-worth-eur-95-million-seized-in-global-operation ↩

[3] CNN — *Italian authorities confiscate almost $1 million in fake olive oil* (July 11, 2024; the Carabinieri raid in Apulia, 42 t of fake extra virgin, 623 L of chlorophyll, 7 charged, forged labels and stamps). https://www.cnn.com/2024/07/11/europe/italy-olive-oil-seizure-scli-intl/ ↩

[4] Food Safety News — *Spanish and Italian investigators uncover olive oil fraud* (December 2023; Operation Omegabad: 260,000 L seized, 11 arrests, lampante sold as extra virgin, Ciudad Real/Sicily/Tuscany, €91,000 in cash). https://www.foodsafetynews.com/2023/12/spanish-and-italian-investigators-uncover-olive-oil-fraud/ ↩

[5] European Commission (DG SANTE) — *EU coordinated action on honey ('From the Hives')* (2021-2022 samples; 46% of 320 consignments suspected of adulteration with sugar syrup; methodology and participating countries). https://food.ec.europa.eu/food-safety/eu-agri-food-fraud-network/eu-coordinated-actions/honey-2021-2022_en ↩

[6] FoodNavigator — *Geographical Indications worth a mouth-watering €75 billion, study finds* (the value of EU GIs: €74.8 billion in 2017; a price premium of about 2× for protected products). https://www.foodnavigator.com/Article/2020/04/22/Geographical-Indications-worth-a-mouth-watering-75-billion-study/ ↩

[7] European Commission — *EU Agri-Food Fraud Network: reports and publications* (fraud affects ~1% of the industry and costs $10,000-40,000 million/year; 758 suspicions in 2023, +26%; meat, honey and oil at the top). https://food.ec.europa.eu/food-safety/eu-agri-food-fraud-network/reports-and-publications_en ↩

[8] EUR-Lex — *Regulation (EU) 2024/1143 on geographical indications* (in force May 13, 2024; a single PDO/PGI/GI and TSG framework; protection as an ingredient and online). https://eur-lex.europa.eu/eli/reg/2024/1143/oj/eng ↩

[9] Food Safety News — *Four sentenced in France for roles in horse meat scandal* (the 2019 trial; Poujol (Spanghero) 2 years + a €100,000 fine, Fasen 2 years + a permanent ban; Comigel, 500+ t to 13 countries). https://www.foodsafetynews.com/2019/04/four-sentenced-in-france-for-roles-in-horse-meat-scandal/ ↩

[10] FDA — *Economically Motivated Adulteration (Food Fraud)* (the EMA definition, e.g. cheap oil sold as 100% olive oil; the FSMA's preventive controls duty). https://www.fda.gov/food/compliance-enforcement-food/economically-motivated-adulteration-food-fraud ↩

[11] Sensors (MDPI) — *Blockchain and IoT for food traceability in olive oil supply chains* (2024; the state of the art in sealing/anchoring oil provenance and in NMR/IRMS analytical authentication). https://www.mdpi.com/1424-8220/24/24/8189

Who Made It, and Under What Conditions: Due Diligence in the Workforce

"Consider if this is a man / Who works in the mud / Who does not know peace / Who fights for a scrap of bread / Who dies because of a yes or a no."

— Primo Levi, "Shemà," If This Is a Man (1947)

In late January and through February 2024, U.S. customs detained thousands of **Volkswagen Group** luxury cars — **Porsche, Bentley and Audi** models — at American ports. The reason was not the engine, nor the bodywork, nor the German steel: it was a **tiny electronic part**, hidden deep in the chain, with material traceable to **Xinjiang**, the region where the *Uyghur Forced Labor Prevention Act* presumes every good is made with forced labor. Volkswagen declared it had not known the component's northwestern-China origin until a supplier alerted it in mid-January 2024, after the American authorities had already flagged it in December 2023; it promised to replace the part in all cars sold worldwide and worked to free the vehicles [4]. The episode is the labor dimension in its purest state: the cars and their steel were "from Germany," but the question that stopped them at the dock was **who had made one minuscule part far down the chain, and under what conditions**.

This chapter is distinct from chapter 29 — supply chain integrity — which seals **where** a good comes from and its physical custody: the object's origin. Here we speak of the **social** layer superimposed on it: who made it and under what labor conditions, the human rights due diligence record and the workers' testimony, not the material's geography. The chapter's question is: how do you fix, with a certain date and in a way that can be neither backdated nor dismissed as a deepfake, **the proof of what we knew and what we did** about the people at the end of the chain. The Volkswagen episode crystallizes it: a multinational can be blind to a fourth- or fifth-tier supplier it does not see.

A scourge of 28 million people and $236 billion

The magnitude is crushing. The **ILO** estimates that **27.6 million people** were in forced labor on any given day of 2021, 2.7 million more than in 2016; **17.3 million** are exploited in the private economy and **migrant workers** run a risk **three times** higher [1]. Forced labor generates **$236 billion** in illegal profits a year — $64 billion more (37%) than in 2014 — with the highest profits in Europe and Central Asia (**$84 billion**) [2]. As we have seen, it is not a problem of faraway countries: it is a global business model, and part of the money is made at home.

The Volkswagen episode did not come out of nowhere. In February 2024, **Human Rights Watch** published the report *Asleep at the Wheel*, documenting carmakers' exposure — VW, BMW, Mercedes-Benz, GM and others — to Uyghur forced labor in the aluminum and components chain [5]. In June 2023, the **ECCHR** had filed a complaint with Germany's **BAFA**, alleging that Volkswagen, BMW and Mercedes-Benz had breached the German supply chain law over their links to Xinjiang. The risk was not invisible: it was documented and denounced before customs acted.

The risk migrates from clothing to cars

For years, forced labor was associated with textiles; no longer. Between June 2022 and July 2025, U.S. customs (**CBP**) detained **16,755 shipments** valued at **$3.69 billion** under the UFLPA; in 2024 alone it detained goods worth **$1.34 billion**, and almost **48%** of that year's detentions ended in denied entry [3]. The figure that says it all: detentions in the **automotive and aerospace** sector grew almost **1,600%** in 2024 over 2023 — from 59 to 991 shipments between January and November — proof that the forced labor risk has migrated from clothing to complex manufactured goods [3].

And the pressure grows. On **January 14, 2025**, the **DHS** made its biggest expansion of the UFLPA *Entity List*, adding **37 Chinese** energy and textile companies and taking the rebuttable-presumption blacklist to **144 entities** [3] — as we have seen. Every added name turns a hitherto quiet chain into a documentary minefield: the buyer must prove, with "clear and convincing" evidence, that their product contains none of them.

The law already demands watching the whole chain

The regulatory framework has hardened on both shores of the Atlantic. In the **U.S.**, the **UFLPA** (in force since June 21, 2022) presumes that any good made wholly or in part in Xinjiang, or by a listed entity, is the product of forced labor and is barred, unless "clear and convincing" evidence rebuts it. In the **EU**, the **Corporate Sustainability Due Diligence Directive (CSDDD, Directive 2024/1760)** entered into force on **July 25, 2024**, covering companies with 5,000 employees and €1.5 billion in turnover (and non-EU ones with €1.5 billion in EU turnover); after the Omnibus amendments, transposition is postponed to **July 26, 2028** and application to **July 26, 2029** [6]. The **European forced labor Regulation**, in force since **December 13, 2024** and applicable from **December 14, 2027**, bans any product made with forced labor, but without a geographic presumption — the burden of proof rests on the authorities, unlike the UFLPA [7]. In **Germany**, the **LkSG** has applied to companies of 1,000 employees since January 1, 2024; the BAFA can impose fines of up to **€800,000 or 2% of worldwide turnover**, and in 2023 it carried out **486 audits** [8]. Beneath it all lies the international framework: ILO Conventions 29 and 105 and the UN Guiding Principles on Business and Human Rights.

Why the social audit alone does not save

It must be said bluntly, because complacency here literally kills. The tool the industry says watches over labor — the social audit and certificates like SA8000, BSCI or SMETA — fails repeatedly at the very thing it exists for. The Rana Plaza factory, audited under BSCI, collapsed on April 24, 2013 and killed 1,131 people; the Ali Enterprises factory in Pakistan, certified SA8000, burned in 2012 and killed more than 250 workers. Both had passed inspection shortly before the disaster [9]. These are not isolated cases: Human Rights Watch's November 2022 report documented how audits are routinely rigged — consultancies that coach workers to lie, double bookkeeping, fake payrolls and children hidden during the inspection [10].

And there is a more fundamental void, prior to the fraud: the certificate only verifies compliance with a legal minimum that is often misery. As we have seen, legal wages sit far below a living wage in all the big textile hubs — around 31% in India, 37% in China and Bangladesh, 33% in Cambodia and 13% in Sri Lanka [11]. A compliance PDF can say "all correct" about a factory where no one earns enough to live. The document, alone, does not save: it is as reliable as the will of whoever signs it and as honest as the inspection that issued it.

The proof of existence applied

This is where **t2p** supplies the missing layer, and it does so on the **contemporaneous, contextualized evidence** of the conditions and the people, not on the product — as we have seen in chapter 29. At the moment of capture, a worker, a union representative, an unannounced field auditor or a trusted local NGO records with their device a photo or a short video of the workshop, of a dormitory, of a chained emergency exit, of a payslip or a shift sheet. The app computes on the device the **SHA-512 fingerprint** of the file — which never leaves the phone — obtains an **RFC 3161/eIDAS timestamp** and anchors the fingerprint on the public blockchain, producing an **.asice** container verifiable offline. For field captures of a person or a place, a **forensic GPS** fix is added (proving the photo was taken inside that factory or that recruitment dormitory) and a **proof of life** (proving it was a real worker, not a coached deepfake witness, who gave the account at that moment).

The same mechanism seals the corporate face of the chain: the **supplier declaration** and the custody attestation a brand collects to rebut the UFLPA's presumption, the **due diligence record** of the LkSG or the CSDDD, a **worker's complaint** filed through a WOVO-style app, or an unannounced auditor's **raw findings** before any factory has had time to "prepare." It complements the existing system by attacking the two holes documented by regulators and HRW: audits are rigged with coached workers and fabricated documents, and a brand fighting a CBP detention must prove **when** it really knew things and when it acted. A complaint or a finding sealed with date, GPS and proof of life is far harder to backdate, edit or dismiss as a deepfake than a PDF report — and this strengthens the worker's testimony, the complaints trail and the file a company presents to the CBP, the BAFA or an EU competent authority.

Now the limit, in all its rawness, because deceiving oneself here would do harm. "Garbage in, garbage on-chain": the seal proves **the act** — that this person, in this verified place and moment, recorded this image or signed this declaration, and that nothing has been altered since. **It does not prove that the**

conditions are good or the wage decent; a sealed, geolocated photo of a tidy workshop staged for the inspection is still a staging. It does not confirm that a supplier's declaration is truthful, or that the third-tier factory shown is really the one making the part, or that a worker's testimony was given freely and not under coercion — the proof of life defeats a synthetic face, not a frightened one. It does not verify that goods are free of Xinjiang content or satisfy the UFLPA or the CSDDD and LkSG obligations; all of that still demands investigation, documentary tracing, unannounced visits and human judgment. The seal makes the evidence unalterable and dated; it never replaces the audit, the regulator or the worker's own credibility. It is the labor complement of what chapters 13 and 15 propose with C2PA and content provenance.

The promise and the peril

The promise is to turn the most fragile testimony in the chain — the voice of a worker, often frightened and powerless — into **dated, geolocated, admissible evidence** that no factory can erase and no lawyer can dismiss as fabricated afterward. The honest brand gains a due diligence file that holds before the CBP and the BAFA; the worker gains that their word stops being "their word against the company's."

The peril calls for the strictest nuances. The first, and it is life-or-death: **the witness's safety**. A worker who seals evidence against their employer faces reprisals; the system must protect their identity, not expose it. The second, do not confuse **integrity with truth**: a false, sealed supplier declaration is still false — and now immutable besides; the seal could breed a false reassurance ("it is anchored, therefore compliant") that corresponds to nothing. And the third, **concentration**: if only the big brands can seal and document, the standard can become a barrier that expels the honest small supplier instead of strengthening the worker. The proof of existence is a tool of dignity and accountability, not a certificate of goodness.

It's time to act

For **workers and unions**: to be able to seal on the device a complaint, a payslip or an image of the workshop as dated, geolocated evidence, without the photo leaving the phone and with identity protected. For **brands and buyers** (Volkswagen, BMW, Mercedes-Benz and the rest): seal the supplier declaration and the due diligence record at the moment of collection, so the file presented to the CBP or the BAFA has a certain date and does not look fabricated afterward. For **field auditors and NGOs**: anchor the raw findings before any factory "prepares," turning the unannounced inspection into unalterable evidence. For **regulators** (the CBP, the BAFA, the EU authorities under the forced labor Regulation): recognize the timestamped sealed evidence as valid proof of what a company knew and when, and treat the certain date as the difference between diligence and alibi. And for **legislators** (the UFLPA, the CSDDD, the LkSG): require the due diligence trail to be verifiable and anchored, not a PDF that can be touched up tomorrow.

At American ports, thousands of luxury cars stopped over a part no one at headquarters knew existed. Giving every worker at the end of the chain a sealed proof of what they lived through, and every honest brand a proof of what it knew and when it acted — never forgetting that the seal does not make the conditions decent — is, on this front, the moment to act: time2act.

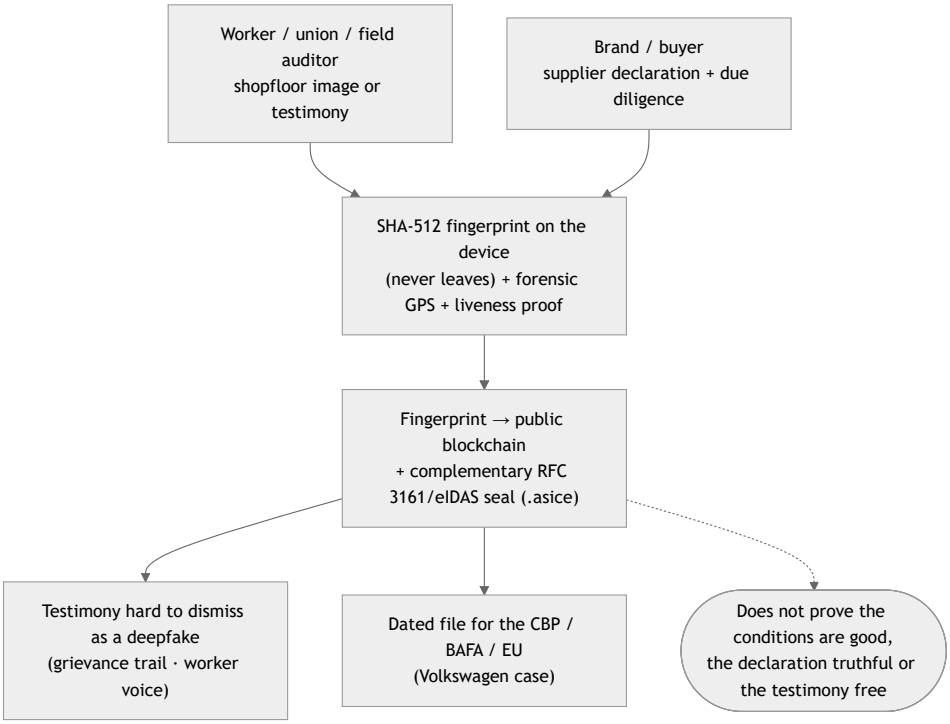


Figure 102.1 — Two sources of evidence: the worker, the union or the field auditor captures on the device an image or a testimony of the workshop (SHA-512 fingerprint that does not leave the phone, forensic GPS, proof of life); the brand seals the supplier declaration and the due diligence record. Everything is sealed with an RFC 3161/eIDAS timestamp anchored on the public blockchain (.asice). Outcomes: worker testimony harder to dismiss as a deepfake, and a dated file for the CBP, the BAFA or the EU authority. Limit: the seal proves the act, the place and the date; it does not prove that the conditions are good, the declaration truthful or the testimony free. Sources: the Volkswagen/CBP case; the ILO; Miller & Chevalier (UFLPA); HRW (audits and automotive); the ECCHR (certificates).

---

[1] ILO — *Data and research on forced labour (Global Estimates 2022)* (27.6 million people in forced labor in 2021; 17.3 million in the private economy; migrants at triple the risk). https://www.ilo.org/topics/forced-labour-modern-slavery-and-trafficking-persons/data-and-research-forced-labour ↩

[2] ILO — *Annual profits from forced labour amount to US$236 billion* (March 2024; +37% since 2014; $84 billion in Europe and Central Asia). https://www.ilo.org/resource/news/annual-profits-forced-labour-amount-us-236-billion-ilo-report-finds ↩

[3] Miller & Chevalier — *UFLPA Enforcement 2024 Year in Review* (16,755 shipments and $3.69 billion detained; $1.34 billion in 2024; 48% denied entry; +1,600% in automotive; the *Entity List* at 144). https://www.millerchevalier.com/publication/trade-compliance-flash-uflpa-enforcement-2024-year-review ↩

[4] Business & Human Rights Resource Centre — *Volkswagen cars held up at US ports over part linked to allegations of forced labour in Xinjiang* (the brands, the sub-component, the December 2023 – mid-January 2024 chronology). https://www.business-humanrights.org/en/latest-news/volkswagen-cars-held-up-at-us-ports-over-part-linked-to-allegations-of-forced-labour-in-xinjiang/ ↩

[5] Human Rights Watch — *Asleep at the Wheel: Car Companies' Complicity in Forced Labor in China* (February 2024; VW, BMW, Mercedes, GM and the aluminum chain). https://www.hrw.org/report/2024/02/01/asleep-wheel/car-companies-complicity-forced-labor-china ↩

[6] European Commission — *Corporate Sustainability Due Diligence (CSDDD)* (Directive 2024/1760, in force July 25, 2024; transposition by July 26, 2028, application from July 26, 2029 after the Omnibus). https://commission.europa.eu/business-economy-euro/doing-business-eu/sustainability-due-diligence-responsible-business/corporate-sustainability-due-diligence_en ↩

[7] Covington — *EU Forced Labour Regulation Enters Into Force* (December 2024; in force December 13, 2024, application December 14, 2027; an investigation model without a geographic presumption). https://www.cov.com/en/news-and-insights/insights/2024/12/eu-forced-labour-regulation-enters-into-force ↩

[8] BAFA — *Supply Chain Act (LkSG): Overview* (application to companies of 1,000 employees since January 1, 2024; fines up to 2% of worldwide turnover; 486 audits in 2023). https://www.bafa.de/EN/Supply_Chain_Act/Overview/overview_node.html ↩

[9] ECCHR — *More for show than safety: Certificates in the textile industry* (Rana Plaza, 1,131 dead, BSCI; Ali Enterprises, more than 250 dead, SA8000; both had passed inspection). https://www.ecchr.eu/en/case/more-for-show-than-safety-

certificates-in-the-textile-industry/ ↩

[10] Human Rights Watch — *Obsessed with Audit Tools, Missing the Goal: Why Social Audits Can't Fix Labor Rights Abuses* (November 2022; coached workers, double books, fake payrolls, hidden children). https://www.hrw.org/report/2022/11/15/obsessed-audit-tools-missing-goal/why-social-audits-cant-fix-labor-rights-abuses ↩

[11] COSH — *Living wages for EU garment workers* (the legal minimum wage is 31% of a living wage in India, 37% in China and Bangladesh, 33% in Cambodia, 13% in Sri Lanka). https://cosh.eco/en/articles/living-wages-eu-garment-workers ↩

Maritime Integrity: Where Is That Ship Really, and Is It Who It Claims to Be?

"The sea has never been friendly to man. At most it has been the accomplice of human restlessness."

— Joseph Conrad, The Mirror of the Sea (1906)

On April 29, 2026, the maritime intelligence firm **Windward** uncovered an Iranian crude-smuggling scheme worth some **$800 million**. Four U.S.-sanctioned VLCC supertankers — **Alicia** (IMO 9281695), **RHN** (IMO 9208215), **Star Forest** (IMO 9237632) and **Aqua** (IMO 9248473) — each with some **2 million barrels** aboard, were broadcasting fake AIS destination messages and erratic tracks to appear anchored off Basra and loading at the Iraqi port of **Khor al Zubair**. In reality, they slipped away to Iranian terminals, loaded Iranian crude and manufactured a **digital alibi** of legitimate Iraqi origin. Windward tracked a dozen sanctioned tankers faking positions off Iraq, and more than twenty vessels west of the Strait of Hormuz with telltale spoofing signs: circular "crop-circle" tracks, invented port calls [2].

This chapter is distinct from chapter 54 — the documentary fraud of international trade, which seals delivery notes and letters of credit — and from chapter 72 — illegal fishing, which asks whether a catch was made where and how it should have been. Here the sealed object is no paper and no catch: it is the **physical fact of where a ship is and who it is**, upstream of any document. At sea, the declared position and identity are **self-reported radio assertions** that can be fabricated wholesale. The chapter's question is how to let a vessel **prove non-repudiably** where it was and who it was at a given instant — and, above all, how to make sure a sanctioned vessel **cannot** fabricate after the fact a trail placing it at Khor al Zubair when it was never there.

The signal that can be invented

AIS — the automatic identification system — was born for safety: so ships can see each other. But what a vessel broadcasts (position, course, MMSI identity, destination) is what **it itself decides** to broadcast. Therein lies the fraud. Between January 2024 and July 2025, **Kpler** identified **261 vessels** that faked their AIS before being sanctioned — the most numerous category of deceptive practice — including small gas carriers inventing calls at Khor al Zubair to fabricate plausible voyage histories [3]. The British firm **Lloyd's List Intelligence** documents the mechanics of *identity laundering*: IMO number fraud, flag hopping, manipulating the AIS "handshake" so two hulls swap identities at sea [11]. The other face is deliberate switching-off: in December 2024, OFAC designated the operators of the **MS ENOLA**, which "continued to employ deceptive practices, including turning off the AIS" to hide illicit transfers of Iranian oil under executive order 13902 [9].

The shadow fleet and the map that lies

The scale is industrial. The **Russian shadow fleet** counts more than **1,000 aging, underinsured tankers** — around **17%** of the world's oil fleet; the Ukrainian intelligence catalogue listed **1,337** in February 2026, and in 2024 the fleet generated some **$9.4 billion** in additional Russian revenue by dodging the price cap [6]. They live on falsification: invented positions, borrowed identities, ship-to-ship (STS) transfers far from any gaze. And it is not only oil: on December 25, 2024 the tanker **Eagle S**, flying the Cook Islands flag and part of the shadow fleet, dragged its anchor and damaged the Estlink 2 submarine cable between Finland and Estonia, adding a **hybrid security** dimension [10].

Above the deliberate fraud sits the **background noise** that makes it indistinguishable: massive GNSS interference. In the second quarter of 2025, more than **10,000 vessels** officially suffered GPS interference — eight times the roughly 1,250 of the first quarter — with almost **3,000 vessels** affected in the Persian Gulf and Hormuz during the June 2025 peak [4]. During that month's Israel-Iran war, spoofing in Hormuz — through which **20%** of the world's oil passes — put supertankers' screens "on dry land" and through a nuclear power plant; Professor Todd Humphreys, of the University of Texas, said it plainly: "I'm quite sure it's the Iranians doing this spoofing" [1]. On March 1, 2026, more than **1,100 Gulf vessels** suffered interference within 24 hours, with 21 new hotspots placing ships inside the Barakah nuclear plant; Hormuz transits fell by a third [5]. The global tracking map has itself become a **forgery**.

The law demands traceability; the fraud ridicules it

The framework exists. The IMO's **SOLAS Convention, Chapter V, Regulation 19** requires AIS on all ships of 300 GT or more on international voyages and requires keeping it "in operation at all times" unless safety justifies switching it off — which turns "dark" sailing and deliberate spoofing into a **safety violation**. **OFAC** sanctions, under order 13902, and the Russian price cap regime explicitly cite AIS manipulation, dark sailing and identity laundering as deceptive practices [9]. The EU had listed more than **444 shadow fleet vessels** by December 2025 (adding 41 on December 18) [7], and the 18th package, of July 18, 2025, lowered the Russian crude cap from $60 to **$47.60** with a dynamic mechanism keeping it 15% below the market [8]. The IMO's Legal Committee is working against **fraudulent registration**: LEG 111 (April 2024) and LEG 112 (March 2025) produced ship registration guidelines, expected for 2027, to curb flag hopping and IMO number fraud. The law demands traceability; the problem is that it rests on a signal the offender itself controls.

Why the AIS mandate alone does not save

It must be said bluntly. That SOLAS requires the AIS to be on prevents neither the switch-off nor the spoofing; it only makes them, on paper, a violation. The underlying problem is that AIS is **self-reported**: the vessel states where it is and who it is, and the receiver believes it. An AIS position without an independent anchor does not distinguish the honest ship from the one fabricating a track; an AIS outage can be a breakdown or a cover; and when the fraud hides inside a sea of **numerous** vessels under real interference, the malicious maneuver becomes indistinguishable from legitimate noise. Detecting the anomaly — which Windward, Kpler or Pole Star do very well — is not the same as **proving it in court**, nor does it stop a sanctioned operator from claiming its suspicious trail was just another victim of the interference. What is missing is a layer that turns "this track looks anomalous" into "this verified position is cryptographically provable, and the absence of a seal is itself a red flag."

The proof of existence applied

This is where **t2p** closes the integrity hole of self-reported position and identity. When a fix is taken, a tamper-resistant onboard unit — or a trusted inspector from the coast guard or the port, or a port-call or bunkering application — captures the **navigation record**: the GNSS fix with time, IMO and MMSI, the draft and the AIS frame as broadcast. On the device itself the record's **SHA-512 fingerprint** is computed — the raw trail need not leave the ship — an **RFC 3161 timestamp** conforming to eIDAS is added and the fingerprint is anchored on the public blockchain, inside an **.asice** container verifiable offline. When an inspector physically checks the hull, the name and the IMO plate at a port call or an STS transfer, a **forensic GPS fix** and a **proof of life** of the inspector and the scene are incorporated.

The result is a **position-and-identity ledger**, continuous and independently anchored: a vessel can later prove where and who it was, and — as we have seen — a sanctioned vessel **cannot** fabricate after the fact a sealed trail placing it at Khor al Zubair when no contemporaneous seal exists. This **complements** the existing systems (AIS, LRIT, satellite RF and SAR detection, OFAC and EU screening) by adding a dated, non-repudiable, court-grade anchor. As chapters 13 and 15 do with the C2PA provenance of what a device **produces**, here what a vessel **does and is** gets sealed.

The seal proves the origin, the time, the author and the integrity of a position record — **not that the underlying GNSS signal was genuine**. If the receiver takes in a spoofed satellite signal, a valid seal will anchor a false position: the cryptography certifies "this device recorded this fix at this instant," not that the fix was real. So without multi-constellation receivers and forensic analysis, the seal does not distinguish an honest location from a spoofed one; only the absence of a seal becomes evidence. It complements C2PA/SynthID provenance but **never replaces** flag audits, port state control or the human judgment that knows how to read a sea.

The promise and the peril

The promise is to turn the ocean into a space with **verifiable memory**: for an honest vessel to be able to prove its position and identity with an anchor no operator can retouch, and for silence — the lack of a seal — to become incriminating data. In a sector where a hull changes name, flag and IMO number without

leaving a receipt, making the origin unalterable protects the clean shipping company, the insurer and the coastal state alike.

The peril calls for the usual nuances. The first, do not confuse **integrity with truth**: a spoofed but sealed position is still false — and now immutable besides; the seal could breed a false reassurance ("it is anchored, therefore real") that corresponds to nothing. The second, **voluntariness**: no regime can force a ghost tanker to seal its position, so the system protects mostly the honest and leaves the proof in the absence. And the third, the **governance** of whoever operates the onboard units and the keys: if they are controlled by whoever has incentives to lie, the seal loses its meaning. The proof of existence is a signal integrity layer, not a certificate of legality or a substitute for the human eye that knows how to read a sea.

It's time to act

For **flag states and the IMO**: build the position-and-identity seal into the 2027 registration guidelines and the SOLAS regime, so that an AIS gap without a contemporaneous seal weighs as evidence. For **port state control authorities and coast guards**: seal with forensic GPS and proof of life every inspection of hull, name and IMO plate at port calls and STS transfers. For **sanctions regulators** (OFAC, the Council of the EU): recognize the absence of a sealed trail as a risk signal in screening, and the seal as court-grade evidence. For **insurers and P&I clubs**: require the anchored position ledger as a condition of coverage, so the shadow fleet's underinsured tanker is pushed out of the market. And for the **maritime intelligence firms** (Windward, Kpler, Pole Star): connect their anomaly detection to the sealed anchor, so that "looks like spoofing" becomes "provably was not where it claimed."

As we saw on April 29, 2026, four supertankers painted onto the map a geography that did not exist to hide $800 million of Iranian crude. Giving every honest vessel a sealed proof of where it was and who it was — and turning the fraudster's silence into his own accusation — never forgetting that the seal does not make a manipulated signal real, is, on this front where the ocean itself has become forgeable, time2act.

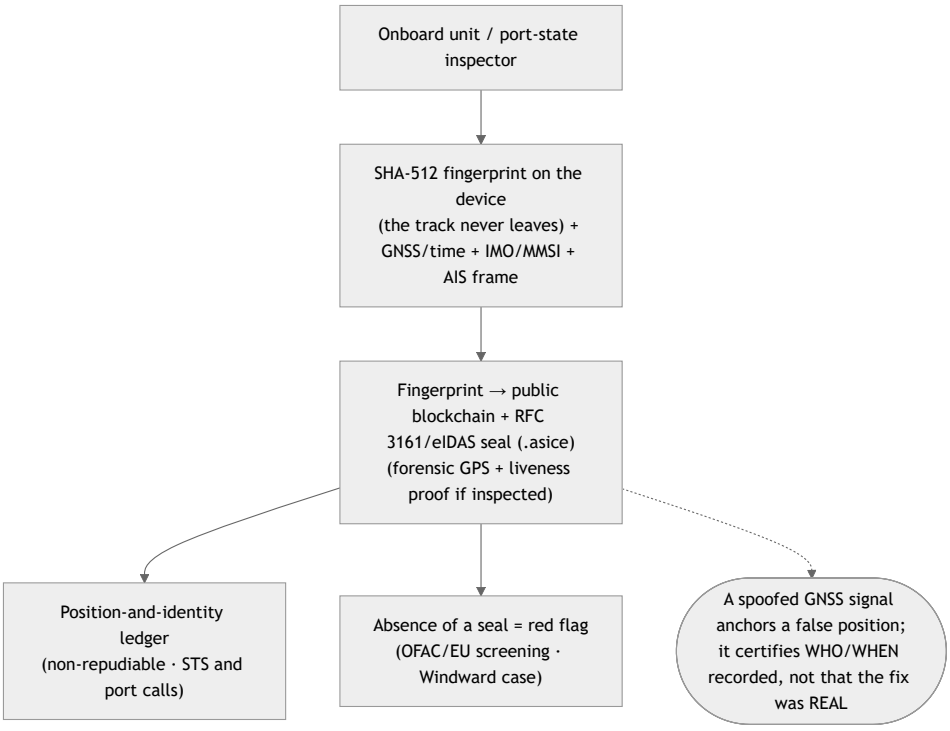


Figure 103.1 — When a fix is taken, an onboard unit or a port state inspector captures the navigation record (GNSS + time, IMO/MMSI, draft, AIS frame) and computes its SHA-512 fingerprint on the device; the whole's fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), with forensic GPS and proof of life when there is a physical inspection. Two outcomes: on the left, a non-repudiable position-and-identity ledger; on the right, the absence of a seal becomes a red flag for OFAC/EU screening. Limit: the seal certifies that the device recorded this fix at this instant, not that the GNSS signal was genuine; a spoofed signal anchors a false position. Sources:

the Windward case (the $800 million scheme); Kpler (261 vessels); GPSPATRON (10,000+ vessels); Scientific American (Hormuz); OFAC (MS ENOLA).

---

[1] GPS spoofing is scrambling ships in the Strait of Hormuz — *Scientific American* (spoofing in Hormuz during the 2025 Israel-Iran war; ships on land and inside nuclear plants; Humphreys's attribution to Iran; ~20% of the world's oil through the strait). https://www.scientificamerican.com/article/gps-spoofing-is-scrambling-ships-in-the-strait-of-hormuz/ ↩

[2] Iran's $800M oil-smuggling scheme uses tankers posing as Iraqi ships — *Fox News* (Windward's disclosure, April 29, 2026; VLCCs with IMO numbers; fake calls at Khor al Zubair). https://www.foxnews.com/world/irans-800m-oil-smuggling-scheme-uses-tankers-posing-iraqi-ships-dodge-blockade-.amp ↩

[3] AIS spoofing: the fast track to sanctions — *Kpler* (261 vessels faked their AIS before being sanctioned, January 2024-July 2025; fake calls at Khor al Zubair). https://www.kpler.com/blog/ais-spoofing-fast-track-to-sanctions ↩

[4] Maritime GNSS Interference Worldwide: A Cumulative Analysis 2025 — *GPSPATRON* (10,000+ vessels affected in Q2 2025, eight times more than Q1; regional hotspots). https://gpspatron.com/maritime-gnss-interference-worldwide-a-cumulative-analysis-2025/ ↩

[5] Widespread GPS jamming hits 1,000-plus ships in the Middle East Gulf — *Windward* (March 1, 2026: 1,100+ vessels in 24 h, 21+ hotspots; false fixes at Barakah and Al Hamra; Hormuz down a third). https://windward.ai/blog/gps-jamming-disrupts-1100-ships-in-the-middle-east-gulf/ ↩

[6] Russian Shadow Fleet Tracker — *Kyiv School of Economics (KSE) Institute* (the fleet's size; ~17% of the world's oil fleet; $9.4 billion in extra revenue in 2024; the growing weight of the Russian flag). https://kse.ua/about-the-school/news/russian-shadow-fleet-tracker-march-2026-share-of-russian-flagged-shadow-tankers-jumps-from-3-to-21-in-nine-months/ ↩

[7] Council sanctions 41 vessels of the Russian shadow fleet — *Council of the EU (Consilium)* (December 18, 2025: the EU total reaches 444+ vessels; port access and service bans). https://www.consilium.europa.eu/en/press/press-releases/2025/12/18/russia-s-war-of-aggression-against-ukraine-council-sanctions-41-vessels-of-the-russian-shadow-fleet/ ↩

[8] EU adopts 18th package of economic and individual measures — *Council of the EU (Consilium)* (July 18, 2025: the crude cap lowered from $60 to $47.60 with a dynamic mechanism 15% below the market). https://www.consilium.europa.eu/en/press/press-releases/2025/07/18/russia-s-war-of-aggression-against-ukraine-eu-adopts-18th-package-of-economic-and-individual-measures/ ↩

[9] Treasury Maintains Pressure on Iranian Shadow Fleet (MS ENOLA) — *U.S. Department of the Treasury / Iran Primer* (December 2024; OFAC action under E.O. 13902 citing the MS ENOLA "turning off the AIS"). https://iranprimer.usip.org/blog/2024/dec/19/us-sanctions-more-irans-dark-fleet ↩

[10] Estlink 2 incident / Eagle S shadow-fleet tanker — *NPR* (the Cook Islands-flagged tanker Eagle S, December 25, 2024, damage to the Baltic submarine cable; the hybrid threat dimension). https://www.npr.org/2024/12/31/nx-s1-5243302/finland-russia-severed-undersea-cable-shadow-fleet ↩

[11] AIS spoofing: vessel-identity manipulation typologies — *Lloyd's List Intelligence* (identity laundering, IMO number fraud, the AIS handshake and the mechanics of position spoofing). https://www.lloydslistintelligence.com/resources/blog/the-secret-lives-of-the-shadow-fleet-ais-spoofing-vessel-identity-manipulation ↩

Earth Observation: The Orbital Witness That Can Also Be Forged

"Look again at that dot. That's here. That's home. That's us."

— Carl Sagan, Pale Blue Dot (1994)

In April 2022, after Russian forces withdrew from **Bucha**, in Ukraine, the Kremlin claimed that the images of bodies strewn across the streets had been "staged" once its soldiers had left. But **Maxar Technologies**' **WorldView** satellites had already photographed those streets: high-resolution imagery from March 18, 19 and 31, 2022 showed corpses lying on Yablonska Street and a freshly dug mass grave — a trench of some 14 meters beside a gold-domed church — all dated to weeks when Russian troops still occupied the town. *The New York Times* geolocated the bodies in the orbital frames to the exact positions later seen in street-level video, demonstrably refuting Moscow's denial and feeding the war crimes investigations [2][3]. And yet that same "witness" in orbit is today a forgery target: in 2021, the geographer **Bo Zhao**, of the University of Washington, coined the term "deepfake geography" and showed, in *Cartography and Geographic Information Science*, that AI could graft Beijing's textures onto a map of Tacoma to fabricate convincing fake satellite scenes [1].

This chapter is transversal and touches three neighbors without merging with them. Chapter 20 uses satellite imagery as **judicial proof** of a war crime; chapter 22 uses it to **measure carbon** (biomass, deforestation); chapter 96 to **monitor emissions**. All three treat the image as a *source* of proof about something else. Here the subject is the image itself: the **integrity of the orbital frame**, at once witness and forgery target. The question is uncomfortable: if the eye in the sky can now lie as well as any deepfake, how do we fix that this pixel, from this sensor, existed at this instant — before anyone disputed it? It is the book's thesis applied to the most unexpected place: the photograph of the Earth from above.

When the orbital image is the only witness

Satellite imagery has become front-line evidence precisely where no one reaches on the ground. In 2020, the **BuzzFeed News** team — the *Built to Last* series, later a Pulitzer winner — combined **Planet** imagery and other satellites with the blank tiles Chinese censorship left on its own maps to identify **268 newly built detention camps** in Xinjiang, and **428 suspect sites** in total [4]. Without a single journalist on the ground, orbit documented a system of mass internment. In Bucha, the pattern repeats in reverse: orbit had dated the crime weeks before the denier could invent an alibi. In both cases, the image was not an illustration: it was *the evidence*, and the whole chain — court, United Nations, public opinion — rested on it.

This dependence is, at the same time, its vulnerability. When an entire edifice of proof rests on a frame, whoever can forge it — or merely sow the doubt that it can be forged — brings the whole edifice down. And this is where 2021 changed the rules.

The witness that can now be manufactured

Bo Zhao did not speak hypothetically. His study built a detection set of **8,064 satellite images** — half authentic, of Seattle, Tacoma and Beijing; the other half deepfakes of Tacoma with textures grafted from Beijing and Seattle — and warned that AI-forged imagery is a **national security threat** [1]. Four years later, the danger is operational, not theoretical. In November 2025, a detection study tested classifiers against aerial imagery forged with **StyleGAN2** and **Stable Diffusion** across more than **130,000 labeled images**, and found that Vision Transformers detect them far better than convolutional networks: **95.11% versus 87.02%** accuracy; GAN forgeries leave high-frequency distortions, while diffusion ones show "semantic dissonance" — repeated terrain, roads that end abruptly, mismatched shadows [5].

The risk rises when the target is strategic. In 2024, nuclear verification researchers (Nguyen, Glaser and Biessmann, of Princeton) showed that text-to-image models can generate **photorealistic satellite images of specific facilities** — nuclear plants included — "difficult to distinguish from real data," with direct implications for arms control and verification [6]. The **NGA** itself — the U.S. geospatial intelligence agency — has warned that GAN-manipulated imagery could divert convoys toward bridges or roads that

do not exist. The eye in the sky, for decades the gold standard of incontestable proof, has entered the era in which it, too, must be asked, "how do I know you are real?"

The institutional race: authenticating the signal and the image

The institutional response advances on two fronts. The first is geolocation: on July 24, 2025 Galileo's OSNMA service entered operation, becoming the world's first GNSS to broadcast a cryptographic digital signature (the TESLA protocol) so receivers can reject spoofed positioning signals [7]. It is the backbone of the where and the when of a capture — though it does not stop jamming. The second front is automatic recognition: as we have seen, in 2024 the NGA launched the Sequoia data-labeling program, worth up to $708 million over seven years — its largest data contract ever — to train computer vision (Maven) on satellite imagery [8]. The more automatic decisions depend on the image, the more expensive it is for the image to lie.

The third front is the trust standard. The EU has stood up its Space Programme (Regulation 2021/696, €14.8 billion for 2021-2027) and the EUSPA agency, charged with the security of Copernicus, Galileo and EGNOS. The Open Geospatial Consortium published in April 2025 an annual Integrity, Provenance and Trust (IPT) initiative that defines integrity (handling across the life cycle), provenance (the origin and who has modified the data) and trust (well documented, consistent, unalterable) for geospatial data [9]. It is the conceptual framework that was missing — but it has a decisive void.

Why the IPT framework alone does not save

It must be said bluntly. The OGC's IPT initiative magnificently defines what we mean by integrity and provenance, but it is based on metadata standards, not cryptographic signing or blockchain [9]. A metadata field is only as reliable as the system that holds it: a provenance field in a file can be rewritten, backdated or softened after a dispute, leaving no independent trace. The same goes for the other layers: C2PA — the Content Credentials that the joint NSA, ODNI and CISA guidance of January 29, 2025 recommends for image integrity — signs provenance if applied at the moment of capture, but says nothing if the forger seals his own forgery. And the Vision Transformer detectors, good as they are, give a probability, not a dated proof: their 95% accuracy means one in twenty forgeries gets through, and every new diffusion generation forces them to retrain. None of these layers fixes, unalterably and independently, who delivered which pixels and when. The chain of custody is missing.

The proof of existence applied

This is where **t2p** supplies the missing layer, and it does so on the digital artifact — the scene — at the earliest defensible point of the chain. The sealed artifact is **the raw or L1-level satellite scene** plus its carrying metadata (sensor identifier, orbit/ephemerides, ground station reception time, projection and processing parameters). There are two complementary anchors. **First**, at ground station ingest: the operator (Maxar, Planet, ESA's Copernicus ground segment) computes the **SHA-512 fingerprint** of the delivered scene on the receiving system — the file does not leave it — obtains an **RFC 3161 timestamp** conforming to eIDAS, anchors the fingerprint on the public chain and issues an **.asice** container verifiable offline that travels with every derived tile and report. **Second**, at the analyst's, journalist's or investigator's end — an OSINT team geolocating the Bucha or Xinjiang frames: the specific crop, the annotation layer and the asserted GPS coordinates and acquisition time are sealed again, with **forensic GPS** binding any corroborating field capture and a **proof of life** when a person verifies a location on the ground.

The seal fixes which pixels existed, from which sensor, at which UTC instant, attributed to which operator or analyst. Thus a scene forged or diffusion-generated later cannot pass itself off as the original, and a court — or an MRV carbon auditor — can prove the image entered the record before any dispute. It complements, does not replace: the OGC's IPT provenance, Galileo OSNMA's position authentication, the C2PA Content Credentials of chapters 13 and 15, and the Vision Transformer deepfake detectors — supplying the tamper-proof custody layer all of them lack.

Now the limit, in all its rawness. "Garbage in, garbage on-chain": the seal proves the origin, the exact instant, the authorship and the bit-for-bit integrity of a scene; it does not prove that the scene shows what is claimed, that the sensor was well calibrated, that cloud, shadow, off-nadir angle or resolution do not

mislead, or that the geolocation is correct. If a forger generates a fake nuclear plant with Stable Diffusion — or a doctored deforestation scene — and seals that, the seal will faithfully date a lie: it only fixes that this exact image existed at that moment, signed by that signer, not that it is authentic Earth observation. It cannot retroactively authenticate an image sealed only after manipulation, it does not stop GNSS jamming, and it does not rule on whether a bridge, a mass grave or a razed forest is real. Detecting synthetic pixels still demands forensic classifiers and expert judgment; MRV carbon, ground inventories; war crimes use, on-the-ground corroboration and judicial assessment. The seal strengthens the chain of custody; it never replaces verifying the truth of the content.

The promise and the peril

The promise is an Earth observation record where every scene carries an **unalterable receipt** of its origin, time and author — so that Bucha's denier can no longer plead "staging" against an image cryptographically dated weeks earlier, and the carbon auditor or the prosecutor can prove the evidence predates the dispute. In a world where the eye in the sky can lie, making the origin unalterable protects the victim and the honest verifier alike.

The peril calls for the usual nuances. The first, do not confuse **integrity with truth**: a scene forged and sealed at the moment of forging is still false — and now carries a seal that can breed a false reassurance ("it is anchored, therefore real") that corresponds to nothing. The second, the **concentration of power**: if only a handful of commercial operators and state agencies can seal at the ground station, the standard can consolidate who controls orbital truth and shut out the independent researcher. And the third, **authoritarian** use: the same chain that gives evidentiary weight to the Xinjiang imagery can give weight to surveillance; the image's integrity does not judge its legitimate use. The proof of existence is a custody layer, not a certificate of authenticity or a substitute for the expert judgment that knows how to read a frame.

It's time to act

For **satellite operators** (Maxar, Planet, ESA's Copernicus ground segment): seal the raw or L1 scene at ground station ingest — SHA-512 fingerprint, eIDAS seal, anchoring on the public blockchain, .asice container — so every derived tile carries unalterable provenance from minute zero. For **OSINT researchers, journalists and human rights NGOs** (those who geolocated Bucha and Xinjiang): seal again the crop, the annotation and the asserted coordinates, with forensic GPS and proof of life when there is a field capture. For **standards bodies** (the OGC, within the IPT initiative): add cryptographic signing and unalterable anchoring to the metadata layer, filling the void the IPT leaves today. For **states and agencies** (EUSPA, the NGA, the Galileo OSNMA program): combine position signal authentication with image sealing, so the when and the where and the pixels are defended together. And for **courts and MRV auditors** (the ICC, carbon verifiers): recognize the dated seal as the chain of custody that distinguishes evidence predating the dispute from an image produced afterward — never forgetting that the seal does not decide whether what is seen is real.

In Bucha, the eye in the sky had dated the crime weeks before the denier could invent the alibi. Giving every satellite scene a sealed proof of its origin, time and author — without forgetting that the seal fixes existence, not the truth of what it shows — is, on this front where the orbital witness is at once our best evidence and a new forgery target, time2act.

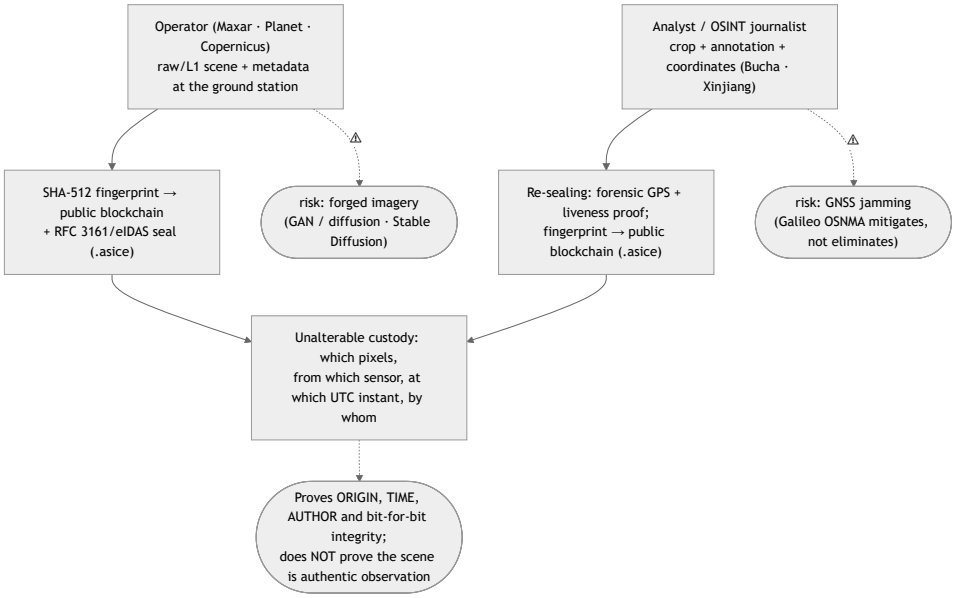


Figure 104.1 — The Earth observation chain in two sealing links. At the origin, the operator (Maxar, Planet, Copernicus) computes the SHA-512 fingerprint of the raw/L1 scene plus metadata at ground station ingest; the fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). At the other end, the OSINT analyst or journalist seals again the crop, the annotation and the coordinates, with forensic GPS and proof of life. Risk points: imagery forged with GAN/diffusion (Stable Diffusion, StyleGAN2) and GNSS interference, which the Galileo OSNMA signature and the Vision Transformer detectors mitigate but do not eliminate. Limit: the seal proves origin, time, author and bit-for-bit integrity; it does not prove the scene authentic observation or the interpretation correct. Sources: the Bucha case (Maxar); Bo Zhao (deepfake geography); Xinjiang (BuzzFeed); arXiv 2511.17766; Galileo OSNMA; OGC IPT.

---

[1] University of Washington — *A growing problem of 'deepfake geography': How AI falsifies satellite images* (April 21, 2021; Bo Zhao coins "deepfake geography" in *Cartography and Geographic Information Science*; the 8,064-image set; a national security threat). https://www.washington.edu/news/2021/04/21/a-growing-problem-of-deepfake-geography-how-ai-falsifies-satellite-images/ ↩

[2] The Moscow Times — *Satellite Images Show Bodies in Bucha for Weeks, Rebutting Moscow Claim* (April 5, 2022; nine Maxar images of March 18/19/31; bodies on Yablonska Street; a 14-meter mass grave). https://www.themoscowtimes.com/2022/04/05/satellite-images-show-bodies-in-bucha-for-weeks-rebutting-moscow-claim-a77211 ↩

[3] Aerospace America / AIAA — *Maxar describes its role in Bucha revelations* (the WorldView-3 capture; Daniel Jablonsky on Maxar's intensive collection; the New York Times's use on April 4, 2022). https://aerospaceamerica.aiaa.org/maxar-bucha-insights/ ↩

[4] BuzzFeed News — *Blanked-Out Spots On China's Maps Helped Us Uncover Xinjiang's Camps* (the Pulitzer-winning method; Planet imagery and censored tiles; 268 new camps and 428 suspect sites). https://www.buzzfeednews.com/article/alison_killing/satellite-images-investigation-xinjiang-detention-camps ↩

[5] *Deepfake Geography: Detecting AI-Generated Satellite Images* (arXiv 2511.17766, November 2025; ViT 95.11% versus CNN 87.02% on more than 130,000 images; GAN versus diffusion artifacts). https://arxiv.org/abs/2511.17766 ↩

[6] *Generating Synthetic Satellite Imagery With Deep-Learning Text-to-Image Models* (arXiv 2404.07754, April 2024; Nguyen, Glaser and Biessmann; photorealistic synthetic imagery of facilities including nuclear plants; arms control). https://arxiv.org/abs/2404.07754 ↩

[7] European Commission — *Observer: How Galileo OSNMA helps counter GNSS spoofing* (OSNMA operational July 24, 2025; the world's first Open Service signal authentication of a GNSS via TESLA cryptographic signatures). https://defence-industry-space.ec.europa.eu/observer-how-galileo-osnma-helps-counter-gnss-spoofing-2025-09-08_en ↩

[8] SpaceNews — *NGA to launch $700 million program to help AI make sense of satellite images* (the Sequoia labeling contract of up to $708 million over 7 years, 2024; feeds the Maven computer vision; the agency's largest data contract). https://spacenews.com/nga-to-launch-700-million-program-to-help-ai-make-sense-of-satellite-images/ ↩

[9] Open Geospatial Consortium — *Navigating Synthetic Imagery Trust in Geospatial Data* (the Integrity, Provenance and Trust initiative, April 2025; defines integrity/provenance/trust, but with no cryptographic signing or blockchain). https://www.ogc.org/blog-article/navigating-synthetic-imagery-trust-geospatial-data/ ↩

[10] NSA/ODNI/CISA — *Content Credentials: Strengthening Multimedia Integrity in the Generative AI Era* (January 29, 2025; joint guidance backing C2PA provenance signing for image integrity, applicable to Earth observation). https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF

Civil Protection and Emergencies: Who Knew What, and When

"But if the watchman see the sword come, and blow not the trumpet, and the people be not warned [...] his blood will I require at the watchman's hand."

— Ezekiel 33.6

On the afternoon of October 29, 2024, the gauging station on the Poyo ravine, southwest of Valencia, sent its last data. At 16:13, an email from the Júcar Hydrographic Confederation (CHJ) to the Emergencies agency reported a flow of 28.7 cubic meters per second: the thing seemed to be easing [5]. At 17:00, the flow was already 325.52; at 17:40, the CHJ itself recorded in the Poyo a flow five times that of the Ebro [6]; at 18:00, the series read 1,725.89 [5]. At 18:43, another email warned: 1,686 cubic meters per second and elevated risk [5]. At 18:55, the sensor read 2,282. Toward 19:00, the force of the water swept away and destroyed the station itself [5]. And the ES-Alert did not sound on phones until 20:11, when the ravines had already taken whole streets [1]. The toll: 229 dead in the province of Valencia and 237 across Spain, more than 2,600 injured, some 141,000 vehicles and 11,242 homes affected, 800 kilometers of damaged roads [4], and estimated losses of €17.8 billion [3].

This chapter tackles civil protection and emergencies: floods, fires, avalanches, the state of the roads and the alerts to the population. The truth of official communications in general is dealt with in chapter 86; the orbital witness of the satellites, in chapter 104; industrial sensors, in chapter 96; the ordinary insurance claim, in chapter 28. The question running through the following pages is another: how do you reconstruct, so that no one can dispute it, the chronology of an emergency? Because after every catastrophe a second battle begins — judicial, political and insurance-related — that always turns on the same thing: who knew what, and when. This book's answer is to seal every datum at the origin with **t2p**: the SHA-512 fingerprint computed on the device or at the source, the RFC 3161/eIDAS timestamp and the anchoring on the public blockchain, inside an .asice container, as we have seen in chapters 13 and 15 [23]. A proof of existence like this is neutral by construction: it can exonerate an administration or incriminate it, shield the citizen before the insurer and disarm hoaxes. It has no side. That is why it works.

The chronology as battlefield

Let us reconstruct the five decisive hours with what is known. The CECOPI — the body that coordinates the emergency — was convened at 15:00 and met at 17:00; past 18:00, the president of the Confederation raised the possibility of sending mass alerts to phones [2]. The first ES-Alert went out at 20:11, and a second at 21:03; in a level 2 emergency, the management and the request to send belong to the regional government [1]. Between the proposal and the sound on the phones, some two hours passed; between the elevated-risk email and the alert, almost an hour and a half. On those minutes a battle of versions has been built ever since, filling front pages, hearings and crossed statements.

And a judicial investigation. Since March 10, 2025, Investigating Court number 3 of Catarroja, with Judge Nuria Ruiz Tobarra at its head, has been investigating former regional minister Salomé Pradas and former regional emergencies secretary Emilio Argüeso for alleged grave omissions in the management of that afternoon [9]. The Valencia Provincial Court upheld Argüeso's status as a suspect; the Valencian High Court of Justice declined to investigate former president Mazón; in March 2026, the case was still "entering its final stretch" [10]. More than five hundred days to fix a chronology of five hours.

The case's central irony deserves a pause. The judge had to order À Punt to hand over "all image and sound recordings" of the CECOPI meeting, and to demand the original videos from the production company — hired, precisely, by the emergencies secretariat itself — in order to "determine the exact time of each video" [11]. A year after the flood, the central instrument of the investigation was a handful of files whose date no one could attest. This is exactly the problem the proof of existence solves: had every minute, every video and every order been sealed at the moment of creation, the chronology would today be a verifiable fact in a matter of hours, not a procedural hypothesis consuming years. And it is worth underlining: the seal prejudges nothing and points at no one in advance. It only fixes the clock — for everyone.

Data that disappears when it matters most

The Poyo sensor died working: according to the CHJ's statement of November 4, the force of the water swept away the measuring systems toward 19:00, after two hours reflecting the flood [5]. So far, the physical fragility. What happened next illustrates a worse fragility: that of the datum itself. In early November, amid the battle of versions, the Poyo data disappeared from the CHJ's public website, and reappeared days later; the press called it "the information blackout" [7]. And later revelations added unease: key ravines that were not part of the monitoring network and a gauge in Massanassa that was not shown publicly [8].

No conspiracy is needed for a critical datum to vanish. A rebooting server, a website "under maintenance" or a frightened official is enough. But the effect is always the same: the public loses the ability to distinguish error from negligence from concealment, and suspicion fills the void. In the middle of a crisis, when trust is the one resource that cannot be improvised, every datum that disappears erodes all the ones that remain.

The structural solution is to anchor the telemetry at the origin. If every reading of the Automatic Hydrological Information System (SAIH) becomes, at birth, a SHA-512 fingerprint with a timestamp anchored on the public blockchain [23], deleting it from the website no longer makes it disappear: its existence and content remain attested forever, and any citizen can verify that the published series matches the sealed one. Even the gaps become provable: what was not measured cannot be dressed up afterward either, because the absence of dated fingerprints is, in itself, a demonstrable fact [8]. Trust stops depending on the agency's good faith and starts depending on mathematics. The Poyo sensor died sending data; the record of what it sent should never be able to die.

The lie during the emergency

In the days after the flood, the flooded parking garage of the Bonaire shopping center, in Aldaia, became the epicenter of a conspiracy theory: "they are hiding the dead from us." Videos of other garages presented as Bonaire's; refrigerated trucks and hearses stripped of context as alleged corpse transfers; a false municipal report speaking of "more than 220 corpses" removed on the first day; false documents pushing the total above 800 victims. When the official investigation ended, the conclusion was categorical: not one fatality in the garage [12]. It did not matter: the machine was already spinning. Newtral debunked more than 60 hoaxes about the DANA in the first month alone [13], with generative AI as the main enabler — plausible images and videos from a text prompt, voices cloned from a few seconds of audio; and that tide of falsehoods came to cause public order problems and to interfere with the emergency crews' work [14].

In an emergency, doubt kills twice: when the citizen does not believe the official alert, and when they believe the hoax. The crisis of truth we described in chapter 10 has here its most lethal version, because reaction time is measured in minutes and there is no margin for artisanal verification.

The proof of existence does not detect lies — it is no hoax detector — but it enables the inverse strategy: that every official piece be born sealed. Drone photos, statements, evacuation orders, casualty counts: if each element carries its fingerprint and date anchored from the origin, any person, outlet or fact-checker can check its authenticity without depending on anyone, and whatever circulates without a seal falls automatically under suspicion. The "liar's dividend" — dismissing inconvenient evidence as fake — is deactivated when the evidence carries a cryptographic anchor: the official who would want to deny a delay can no longer say the record is a fabrication, and the conspiracist waving an apocryphal "report" collides with an ecosystem where authentic documents are known by their fingerprint. Official truth becomes checkable; the lie, orphaned of a seal.

Fires, snow and asphalt: the pattern repeats

In August 2025, Spain lived through its worst wave of forest fires in more than thirty years: close to 400,000 hectares burned by August 19, with a score of major fires declared in two weeks [17][18]. Ourense, with some 150,000 hectares — a seventh of the province — León and Zamora concentrated close to half the burned area; the Uña de Quintana fire razed some 40,000 hectares and the A Rúa fire more than 38,000; in Molezuelas de la Carballeda, more than 700 people had to be evacuated, and the government deployed 500 additional UME soldiers on top of the more than 1,400 already active [17]. When the smoke cleared, the eternal question began, the same as the DANA's but with fire: was that village warned in time? Every evacuation order, every road closure and every deployment report is a dated document that will later be disputed.

Snow repeats the pattern in silence. The Lauegi Center of Aran publishes every day the avalanche danger bulletin from field outings and snowpack tests, complemented by the ATES terrain exposure mapping, and warns that this information "does not replace training, experience and good judgment" [19]; the ICGC covers the Catalan Pyrenees and AEMET the other ranges [20]. When there is an accident, the question — from insurers, families and, where appropriate, courts — is always the same: what danger level was published when the group set out? A bulletin sealed at publication makes the answer indisputable, and protects everyone at once: the forecaster can attest what they published and when; the victim and their family, what information the decision was made on.

And the asphalt. Article 106.2 of the Constitution and Law 40/2015 oblige the administration to compensate damage caused by the functioning of public services: an accident due to the poor state of the road is compensable if there is a causal link and no force majeure intervenes, the Guardia Civil's report is the central piece of evidence, and the six-month administrative silence pushes the citizen into litigation [22]. Maintenance records, inspections and road closures sealed at the origin settle the argument from both sides: the diligent administration can prove it acted; the driver, that the pothole was there. In all three scenarios — fires, snow, asphalt — the constant is the same as in Valencia: the second battle is won or lost according to who can attest their dates.

The proof of existence applied

Imagine an emergency coordination center, next October. Nothing has visibly changed: the same screens, the same calls. But every minute, every SAIH email, every order and every alert is being sealed as it is produced, with no added effort for the operators: the SHA-512 fingerprint is computed at the source, receives the RFC 3161/eIDAS timestamp and is anchored on the public blockchain inside an .asice container [23]. The chain of custody builds itself, while everyone does their job.

1. The record of the decision

The coordinating body of an emergency — Valencia's CECOPI or any other — produces minutes, videos and orders that direct the response. Sealed at the moment of creation, they become a non-repudiable chronology: if the sending of an alert is raised past 18:00 and executed at 20:11 [1][2], the record attests it without court orders to broadcasters and production companies [11], and an investigation that today consumes years [9][10] largely reduces to verifying fingerprints. Neutrality is the keystone: the same seal that can expose an official's slowness attests the diligence of whoever acted well. The authorities can prove they decided in time; the public holds a verifiable basis for demanding accountability when they did not.

2. The sensor that cannot disappear

The SAIH readings — like the Poyo sequence: 28.7, 325.52, 1,725.89, 2,282 cubic meters per second [5] — are anchored at the origin, in batches and in real time. Neither an "information blackout" on the website [7] nor the physical destruction of the station [5] erases the record: the fingerprint attests what was measured and when, and the series' gaps are dated too. The monitoring network becomes auditable from outside [8]: anyone holding the data file can check that it matches, bit for bit, what was sealed that day. The argument over "what data existed" ends before it begins.

3. The citizen's photo

The Insurance Compensation Consortium has paid 212,237 DANA claims worth €4,378 million, with a total estimated cost of some €4.8 billion (March 2026 data) [15][16]. Every file begins with the same fragile material: photos and videos of the submerged car, the flooded ground floor, the lost harvest. A photo sealed on the device at the moment it is taken — the fingerprint with the date, without the image having to leave the phone [23] — speeds up the appraisal, shields the honest citizen from suspicion and leaves no margin for the fraud of whoever recycles photos from another flood: without a seal coherent with the date of the loss, the recycled image does not stand. The scale makes it decisive: 251,249 files at once [16], an administrative avalanche that distinguishes the emergency from the ordinary insurance flow we saw in chapter 28.

4. The bulletin and the road

The avalanche danger bulletin [19][20], the evacuation order for a village threatened by fire [17], the preventive closure of a road or a highway maintenance record [22] share one nature: they are public

documents whose evidentiary value depends on no one being able to dispute their time of publication. Sealed at publication, they protect at once the technician who warned, the driver who claims and the family that wants to know whether the warning arrived in time. The certain date has no side: it belongs to whoever needs it.

What the seal does not do must be said with the same clarity: it does not predict floods, does not put out fires, does not sound any siren; it does not detect hoaxes or prevent their creation. It only makes indisputable what existed and when. It is evidentiary infrastructure, not a substitute for investment in prevention.

The promise and the peril

The promise is an emergency with all its critical data sealed at the origin: sensors, minutes, alerts, orders, photos. The later reconstruction would take hours, not years [9][10][11]; the Consortium could process the avalanches of files on non-repudiable evidence, with shorter appraisals and less fraud [16]; and the hoax would lose force because everything official would be verifiable by anyone [12][13]. The second battle — the battle of versions — would be defused before it began.

The perils demand the same frankness.

The first is **surveillance**. Sealing the reception of alerts or citizens' photos can never become a centralized trail of where each person was. The design must be the book's: the fingerprints are generated and held on each person's device; no content and no location travels to any central registry, and the system cannot serve to deanonymize anyone [23].

The second is the **defensive administration**. If the seal becomes a bureaucratic shield, the official may end up spending more energy on "being covered" than on deciding well. The umbrella culture already exists; cryptography must not amplify it. The seal documents the action: it neither replaces nor excuses it.

The third is **revocability**. The fingerprints of personal damage photos must be invalidatable at the citizen's request. Immutability is for the public datum — the flow, the minutes, the order — not for the victims' private lives, who keep the right to be forgotten.

And the fourth is **techno-solutionism**. No blockchain saves lives. Lives are saved by sirens, rehearsed emergency plans, cleared ravines, a culture of self-protection and, above all, the decision taken in time [2]. The seal only guarantees that, afterward, the truth is indestructible. But that, properly understood, is prevention too: whoever knows their inaction will be attested forever has a powerful incentive to act.

It's time to act

The Valencia DANA left 229 dead [4], a judicial investigation counting its days by the hundreds [10] and a lesson no country should need to learn twice: the truth of an emergency is too important to entrust to the memory of meetings, to websites "under maintenance" and to videos without a time.

For civil protection administrations and 112 centers: seal minutes, orders and alerts at the origin. ES-Alert — born of Article 110 of European Directive 2018/1972, transposed by Law 11/2022, and operational since 2022-2023 [21] — already reaches every phone; make the chronology of the next emergency a verifiable fact from the first minute, and not a years-long investigation [9][10].

For the hydrographic confederations and sensor operators: anchor every SAIH reading in real time, so that no "information blackout" [7] is ever possible again and even the measurement gaps are attested [8].

For insurers and the Consortium: accept and prioritize the photo sealed on the device. Facing 251,249 files at once [16], non-repudiable evidence means faster appraisal, less fraud and more respect for the honest policyholder.

For the media and fact-checkers: demand and check the seal on official emergency material, and treat whatever circulates without an anchor with caution by default [12][13].

For you, citizen: document the damage with the fingerprint on the device, and remember that the seal does not replace the alert: when your phone sounds [1], leave the car and go upstairs.

For legislators: grant full evidentiary value to the sealed emergency record and impose it as the standard of the National Alert Network [21].

The Poyo sensor died at 19:00 sending the truth [5]. Let no emergency truth ever again depend on the survival of a device or the memory of a meeting. On this front, time2act.

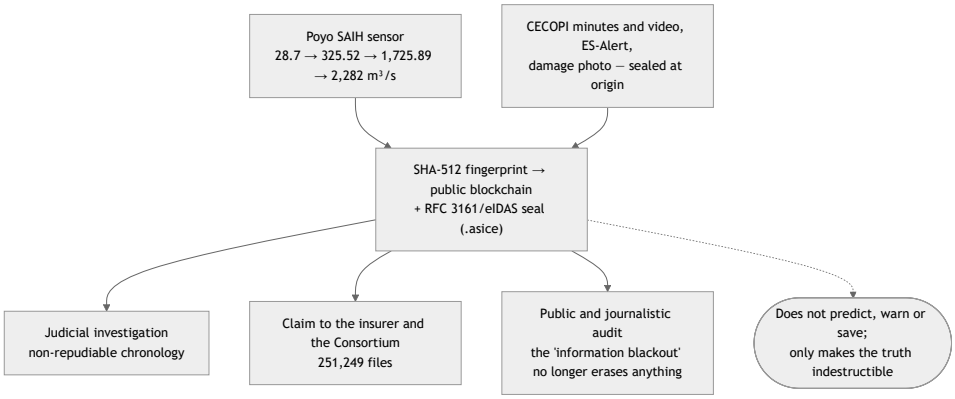


Figure 105.1 — The emergency datum is sealed at the origin: the sensor reading (SAIH/Poyo), the coordinating body's minutes and video (CECOPI), the alert (ES-Alert), the evacuation order and the citizen's damage photo receive the SHA-512 fingerprint on the device or at the source, the RFC 3161/eIDAS timestamp and the anchoring on the public blockchain (.asice). Three outcomes: the judicial investigation (a non-repudiable chronology), the claim to the insurer and the Consortium, and public and journalistic auditing. Limit: the seal does not predict, does not warn and does not save; it only makes the truth indestructible. Sources: Maldita.es; CHJ/SAIH; the Insurance Compensation Consortium; Lauegi.

---

[1] Maldita.es — *Quién envió la alerta de las 20:11 horas en Valencia, qué dice el protocolo* (October 31, 2024; ES-Alert at 20:11 and 21:03; at level 2 the request belongs to the regional government). https://maldita.es/malditateexplica/20241031/generalitat-valenciana-responsable-alerta-movil/ ↩

[2] elDiario.es — *El envío de la alerta masiva a móviles se planteó en la reunión clave de la DANA pasadas las 18.00, dos horas antes* (the CECOPI convened at 15:00, met at 17:00; the mass alert proposed past 18:00). https://www.eldiario.es/comunitat-valenciana/reunion-clave-dana-debatio-enviar-alerta-masiva-moviles-pasadas-18-00-horas_1_11802363.html ↩

[3] France 24 — *En cifras: 229 muertos y 17.800 millones de euros en pérdidas, a un año de la DANA* (October 28, 2025). https://www.france24.com/es/europa/20251028-en-cifras-229-muertos-y-17-800-millones-de-euros-en-p%C3%A9rdidas-a-un-a%C3%B1o-de-la-dana ↩

[4] Wikipedia — *Inundaciones de la DANA de 2024 en España* (237 dead in total: 229 in the province of Valencia, 7 in Castilla-La Mancha, 1 in Málaga; more than 2,600 injured; 11,242 homes; ≈141,000 vehicles; 800 km of roads). https://es.wikipedia.org/wiki/Inundaciones_de_la_DANA_de_2024_en_Espa%C3%B1a ↩

[5] Maldita.es — *Qué sabemos de los datos de la riada en la Rambla del Poyo durante la DANA* (November 7, 2024; the CHJ's emails of 13:42, 15:04, 16:13 and 18:43; the 17:00 and 18:00 flows; 2,282 m³/s at 18:55; the gauging station destroyed by the water toward 19:00, per the CHJ's November 4 statement). https://maldita.es/clima/20241107/datos-rambla-poyo-dana-valencia/ ↩

[6] El Debate — *«A las 17.40 del día de la dana la CHJ registró que el caudal del barranco del Poyo era cinco veces superior al del Ebro»* (February 24, 2026). https://www.eldebate.com/espana/comunidad-valenciana/20260224/1740-dia-dana-chj-registro-caudal-barranco-poyo-cinco-veces-superior-ebro_388667.html ↩

[7] El Debate — *La Confederación del Júcar borra datos críticos del caudal del Poyo en plena crisis de la DANA* (November 5, 2024); esDiario — *Desaparecen y reaparecen los datos del barranco del Poyo de la web de la CHJ*. https://www.eldebate.com/espana/comunidad-valenciana/20241105/confederacion-jucar-borra-datos-criticos-caudal-poyo-plena-crisis-dana-valencia_242165.html · https://www.esdiario.com/comunidad-valenciana/241105/146265/desparecen-aparecen-datos-barranco-poyo-web-chj.html ↩

[8] Actualidad Valencia — *La CHJ omitió barrancos clave en su red de control y ocultó un medidor en Massanassa, según revela el SAIH*. https://actualidadvalencia.com/control-caudal-chj-barrancos-horteta-gallego-massanassa/ ↩

[9] ORM — *La jueza de Catarroja llama a declarar como investigados en la causa por la DANA a Salomé Pradas y Emilio Argüeso* (investigation opened March 10, 2025; Investigating Court No. 3 of Catarroja, Judge Nuria Ruiz Tobarra). https://www.orm.es/informativos/la-jueza-de-catarroja-llama-a-declarar-como-investigados-en-la-causa-por-la-dana-a-salome-pradas-y-emilio-arg-eso/ ↩

[10] El Independiente — *La causa de la DANA afronta su recta final tras el varapalo del TSJ a la jueza de Catarroja* (March 17, 2026; the High Court declines to investigate Mazón). https://www.elindependiente.com/espana/2026/03/17/causa-dana-afronta-recta-final-varapalo-tsj-jueza-catarroja/ ↩

[11] The Objective — *La jueza de la dana une a la causa los vídeos del Cecopi y pide a À Punt la hora de cada uno* (October 2, 2025); El Español — *La jueza requiere a À Punt «todas las grabaciones de imagen y sonido» de la reunión del Cecopi* (the

videos, from a production company hired by the emergencies secretariat; the originals demanded to fix each one's exact time). https://theobjective.com/espana/tribunales/2025-10-02/jueza-dana-videos-cecopi-punt-hora/ · https://www.elespanol.com/valencia/20250919/jueza-dana-requiere-punt-todas-grabaciones-imagen-sonido-reunion-cecopi-octubre/1003743931870_0.amp.html ↩

[12] Maldita.es — *«Nos están ocultando los muertos» del parking de Bonaire (Aldaia): un año después, cómo funcionó la maquinaria de la desinformación* (October 29, 2025; the false municipal report of "more than 220 corpses"; false documents with more than 800 victims; not one fatality in the garage). https://maldita.es/malditobulo/20251029/teoria-conspiracion-muertos-bonaire/ ↩

[13] Newtral — *Un mes y más de 60 bulos desmentidos: la desinformación surgida a raíz de la dana* (November 29, 2024). https://www.newtral.es/desinformacion-dana/20241129/ ↩

[14] COPE — *Quién está detrás de la creación de bulos de la DANA y qué herramientas digitales han usado* (generative AI as the main enabler); elDiario.es — *Los bulos y falsedades sobre la DANA de València* (interference with emergency work). https://www.cope.es/programas/la-linterna/audios/detras-creacion-bulos-dana-valencia-herramientas-digitales-han-usado-20241117_3046424.html · https://www.eldiario.es/comunitat-valenciana/bulos-falsedades-dana-valencia-cientos-cadaveres-bonaire-demolicion-pantanos_1_11791455.html ↩

[15] La Moncloa — *El Consorcio de Compensación de Seguros supera los 4.000 millones en pagos por indemnizaciones a afectados por la DANA* (November 26, 2025; 210,106 payouts). https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/economia-comercio-empresa/paginas/2025/261125-indemnizacion-dana.aspx ↩

[16] SegurosNews — *El Consorcio ya ha pagado 4.378 millones por la DANA de Valencia* (March 27, 2026; 212,237 claims paid; 251,249 files; 99.3% processed; total estimated cost ≈€4.8 billion). https://segurosnews.com/news/el-consorcio-ya-ha-pagado-4-378-millones-por-la-dana-de-valencia ↩

[17] Wikipedia — *Incendios forestales de España de 2025* (≈400,000 ha by August 19; Ourense ≈150,000 ha; Uña de Quintana ≈40,000 ha; A Rúa >38,000 ha; Molezuelas de la Carballeda, more than 700 evacuated; 500 additional UME soldiers on top of more than 1,400). https://es.wikipedia.org/wiki/Incendios_forestales_de_Espa%C3%B1a_de_2025 ↩

[18] Vallfirest — *La oleada de incendios forestales de agosto de 2025* (a score of major fires in two weeks, more than 300,000 ha; the worst season in more than 30 years). https://www.vallfirest.com/noticias/oleada-incendios-forestales-2025 ↩

[19] Lauegi Center of Aran (Conselh Generau d'Aran) — *Avalanche danger bulletin and ATES mapping* (daily bulletin; "the information does not replace training, experience and good judgment"). https://lauegi.conselharan.org/ates/?lang=en ↩

[20] ICGC — *Butlletí de Perill d'Allaus (BPA)* (the Catalan Pyrenees); AEMET — *Boletín de peligro de aludes* (the other ranges). https://www.icgc.cat/en/Thematic-areas/Risks-and-emergencies/Avalanches/Avalanche-Danger-Bulletin-BPA · https://www.aemet.es/en/eltiempo/prediccion/montana/boletin_peligro_aludes ↩

[21] Civil Protection (Spain) — *Public Warning System / ES-Alert* (Art. 110 of Directive (EU) 2018/1972, transposed by Law 11/2022 on Telecommunications; operational since June 2022 and available to the regional 112 centers since 2023; cell broadcast technology). https://www.proteccioncivil.es/coordinacion/redes/ran/public-warning-system · https://www.vademecumlegal.es/legislacion/articulo-110-codigo-europeo-comunicaciones-electronicas ↩

[22] esPúblico — *Responsabilidad patrimonial por daños por el estado de la calzada* (Art. 106.2 of the Constitution; Law 40/2015; the causal link; the Guardia Civil's report as the central evidence; the 6-month administrative silence). https://www.administracionpublica.com/responsabilidad-patrimonial-por-danos-por-el-estado-de-la-calzada/ ↩

[23] t2p — product manual (SHA-512 fingerprint on the device or at the datum's source; RFC 3161/eIDAS timestamp; anchoring on the public blockchain; .asice container). See chapters 13 and 15. ↩

Verifiable Elections: From the "Liar's Dividend" to Anchored Evidence

"If everybody always lies to you, the consequence is not that you believe the lies, but rather that nobody believes anything any longer."

— Hannah Arendt, interview with Roger Errera (1973)

On **September 28, 2023**, two days before Slovakia's parliamentary elections, an audio recording began circulating on Telegram and Facebook in which the leader of the Progressive Slovakia party, **Michal Šimečka**, and the Denník N journalist **Monika Tódová** seemed to be discussing how to rig the elections: buying votes from the Roma minority, manipulating the count and bribing journalists. The audio was an AI-generated forgery; both denied it immediately and AFP's fact-checkers found the signs of manipulation in it. But it arrived at the worst possible moment: inside the **48-hour electoral moratorium**, when neither the media nor politicians could rebut it in detail; and, being audio and not video, it slipped through the hole in Meta's "manipulated media" policy, which then covered only video. It passed **100,000 views** on Facebook and **30,000** on Instagram, and ran through email chains with the transcript included. Šimečka's party lost to Robert Fico's Smer; analysts still debate how much the audio influenced the result, but the case has become the first case study of a lie that arrives when there is no time left to rebut it [12].

Four months later, in New Hampshire, a cloned voice of **Joe Biden** asked voters by phone not to go to the primaries — the case we analyzed in chapter 68 — and the FCC ended up sanctioning the person responsible [1]. The underlying problem is the same on both shores of the Atlantic: when the forgery arrives faster than the denial, electoral trust is left without an anchor. That is this chapter's question.

These scenes become the illustration of a deep paradox: the most fundamental threat is not that we are deceived by a lie, but that we abandon faith in any truth. Beth Noveck, in *Reboot*, describes this phenomenon as the "liar's dividend," where the possibility of creating false narratives allows inconvenient realities to be dismissed as deliberate hoaxes [2]. Thus a simple audio recording can trigger a domino effect that erodes the credibility of the entire democratic process.

The impact of incidents like these is best measured at the scale of the 2024 elections: the biggest electoral year in history, with more than four billion people called to vote in more than forty countries [3]. In a context where every vote counts and a wide fraction of the population depends on a transparent process, the erosion of trust becomes a systemic crisis. Emerging technologies — AI for generating synthetic content and social platforms that amplify its reach — have turned lies into tools of political pressure, creating an environment where the truth needs a more robust defense than the mere detection of what is false.

The answer this chapter proposes, and which forms the connecting thread of Part XII, is not just fine-tuning deepfake filters or reinforcing laws like Utah's SB 131, which requires labeling synthetic media with tamper-evident provenance [4]. The real change of course lies in a proof of existence anchored on the blockchain: the t2p described in chapter 13. This neutral, immutable anchor offers a verifiable record that depends neither on trusted third parties nor on the integrity of the original creator; it is, therefore, an unalterable evidence of authenticity that can be consulted independently of the visual or audio content. The integrity of the process itself — the roll, the ballot box, the count seen from inside — we dealt with in chapter 48; this chapter's front is the public evidence surrounding the election.

Instead of focusing exclusively on detecting the fake — an often losing race, as chapter 15 shows — our mission is to prove what is real with immutable evidence. Only thus can we counter the "liar's dividend" and restore the trust that sustains democratic elections. In this scenario, the blockchain becomes not just a technological infrastructure, but the foundation of a new era of electoral transparency.

The liar's dividend

When technology creates images and audio so perfect that no one can discern the true from the deception, the mere existence of a deepfake becomes a powerful weapon against the truth. Not only does the fake deceive; the **possibility** that any piece of evidence could be an AI-generated fabrication breeds a universal doubt: "maybe all of this is a manipulation." This scenario, called by Beth Noveck the *liar's dividend*, turns every fragment of electoral evidence — from a photograph of the polling station to the

digital record of a vote — into a possible fiction. When no image, audio or tally sheet can be believed with certainty, the vote loses its anchor in shared reality and becomes an isolated gesture, stripped of collective trust.

This erosion of the democratic foundation is no mere hypothesis. As we have seen, 2024 was the biggest electoral year in history, with more than **4 billion** people called to the polls in more than forty countries [3]. In this massive context, any uncertainty about the validity of the results can trigger political crises that cross borders. Deepfake detection, however advanced, becomes an arms race: every new algorithm capable of identifying manipulations is probed and outrun by more sophisticated deception techniques, such as crops or compressions that alter the original signal. As chapter 15 indicates, relying exclusively on detection is a **losing race**; the defenders of the true are constantly outpaced by the evolution of the fake.

The solution is not to chase the false with more zeal. What is needed is a positive foundation, a **neutral, immutable anchor** that can prove, beyond doubt, that something authentic existed and has not been altered. Here enters the blockchain-based proof of existence: the t2p we have already mentioned in chapter 13. This technology offers a cryptographic timestamp verifiable by any actor, without needing a trusted third party or depending on credentials that can be revoked or forged. Instead of hunting the false, the **true** is proven: every electoral tally, every vote record and every live transmission can be "anchored" on the chain, ensuring it has suffered no later alteration.

Watermarks and provenance standards, despite their importance as a complement — as the 2024 *Tech Accord* shows [5] — are insufficient taken separately. Noveck insists that "watermarks and provenance standards alone will not solve the problem"; they must be combined with an immutable infrastructure providing **verifiable trust** at all times [10]. Thus AI stops being an instrument for creating doubt and becomes a tool for reinforcing the truth, tagging every fragment of evidence with an unalterable proof.

In short, the liar's dividend is not just a side effect of deepfakes; it is a structural threat that destabilizes trust in elections. The answer lies not in an endless war against the fakes, but in establishing a neutral, immutable base — the blockchain as the immutable witness of the real fact — that lets every voter trust that their vote is counted as cast, without the shadow of a possible deception. This is the true key to restoring democratic credibility in an era of digital manipulation.

The complement: law and industry (necessary, insufficient)

In the spring of 2024 the state of Utah passed SB 131, a rule requiring any political ad containing "synthetic media" to include a tamper-evident provenance chain: it must be clear who the original author is, which editors intervened and whether generative AI was used. The penalty provided — $1,000 per violation — aims to create an economic disincentive so that political actors cannot hide the truth behind a manipulated image or voice [4]. This law recognizes, explicitly, that the key to fighting deepfakes is tracing the content's origin. Yet its approach rests on a trust model dependent on the creator and the intermediaries: if the author fails to register the information properly, or if a platform decides not to enforce it, the mechanism is switched off.

At the same time, in February 2024 the *Tech Accord to Combat Deceptive Use of AI in Elections* was signed at the Munich Security Conference. Twenty leading companies — Adobe, Amazon, Google, IBM, Meta, Microsoft, OpenAI and others — voluntarily committed to developing tools to detect deceptive electoral content and to "trace the origin" of this material [5]. The accord carries no legal penalties and no external verification mechanism; it rests solely on the participants' good faith. This means that if a company decides to ignore its own guidelines, or if a new provider enters the market without joining the accord, the trust gap persists — as we have seen.

Beth Noveck, author of *Reboot* (2026), puts the final touch on this incomplete picture: "watermarks and provenance standards alone will not solve the problem," because they depend on the cooperation of the creator and the platform [10]. In her words, these solutions are only a "minor complement" to a crisis that requires an independent anchor. AI can help classify and filter content, but without an immutable source that no party can alter, trust remains fragile.

— This is where blockchain anchoring enters as the undisputed protagonist. Unlike the content credentials described in chapter 13 — as we have seen — which rest on an Achilles' heel and depend on a central authority to validate themselves, the proof of existence offers an immutable truth. Once the voter or the electoral record is inscribed on a blockchain with a SHA-512 hash and a Merkle tree, any later attempt to modify it becomes instantly detectable and, more important still, impossible without leaving a trace. This

total neutrality removes the need to trust either whoever created the content or whichever platform hosts it.

— In short, legislative initiatives like Utah's SB 131 — as we have seen — are indispensable steps toward a normative and cultural framework that recognizes the importance of provenance. Yet without an additional layer providing an unalterable proof of existence, these measures fall short. The true solution to the electoral trust crisis is not just tracing the origin; it is creating an immutable witness that any citizen can consult, without depending on third parties. That is what t2p offers: a neutral, independent anchor that turns electoral truth into a public good, immutable and verifiable for all.

The neutral anchor: proving the authentic

In a world where the "liar's dividend" has become a political weapon — "the ability to dismiss true but inconvenient facts as false" [2] — trust in elections stands at the edge of a precipice. The solution is not only to trace the content's origin, but to create an immutable witness that any citizen can consult without depending on any trusted third party; as we have seen, t2p offers a blockchain proof of existence that acts as a neutral, independent anchor, turning electoral evidence into a verifiable public good.

The anchoring rests on three principles. First, the chain takes no sides; each block is generated by a distributed network where no single entity controls the consensus. Second, the creator need keep no credential or secret key: once the evidence is inscribed, its hash becomes an irreversible digital fingerprint anyone can validate against the public chain. Third, this neutrality lets parties, independent observers, media and citizens use the same mechanism without fearing internal manipulation.

A living example of this application is found in the Voto Libre project in Peru and Latin America. There, AI classifies the photographic images collected during the counting phases and, together with blockchain, generates a "secure, transparent, real-time" report [6]. When the Peruvian government passed law 32270 for digital voting with blockchain the following year, it established that every photographic proof of fraud should be anchored on the chain, guaranteeing that no one can alter it without a trace. Thus the images become witnesses verifiable by any citizen or international organization.

In Taiwan, former digital minister Audrey Tang drove a "prebunking" campaign where voters filmed their own polling stations and uploaded that footage to a blockchain-based platform [7]. Each video is registered with a unique hash; when consulted, anyone can check that the content matches exactly the file's original state. This process has turned personal recordings into verifiable proof of what really happened at the polls, countering conspiracy theories and reinforcing the result's legitimacy.

As we have seen in chapter 13, a neutral anchor can be applied to every kind of electoral act: vote counts, official statements from the electoral board, citizens' videos of the polling stations or even candidates' public statements. Once registered on the chain, each element acquires a layer of trust that depends not on the voting device (the "frontend") but on its later record (the "backend"). This means that even if the electronic voting infrastructure is vulnerable, the proof of existence keeps the evidence's integrity intact and lets anyone revisit the process with an immutable truth.

Beth Noveck warns that watermarks and provenance standards "alone will not solve the problem" [10]; they must be combined with mechanisms like the proof of existence to create an ecosystem where media literacy and solid institutions can support real trust. In this sense, t2p offers the missing key piece: a neutral chain that turns every fragment of electoral truth into a public witness, immutable and verifiable for all. (See chapter 13 for the technical detail of the proof of existence.)

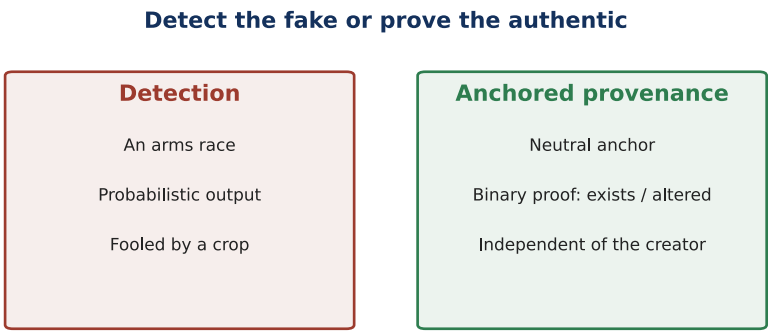


Figure 106.1 — Against the often losing race of detection, the neutral anchor proves that something authentic existed and has not been altered.

When the candidate is an avatar

In the spring of 2024, while votes were being gathered in Wyoming, an unexpected figure appeared before the city council: **VIC**, the Virtual Integrated Citizen. The chatbot, programmed as its creator's "meat puppet," ran for mayor and, surprisingly, won **327 votes** [8]. In American Samoa, another avatar — **PalmerAI**, with the voice and image of Jason Palmer — was proclaimed winner of the Democratic caucus with a **51-40** vote [8]. And in Albania, September 2025 marked the arrival of **Diella**, the world's first AI minister, in charge of public procurement and authorized by a final human sign-off [8].

These cases make the traditional question vanish: "Is this video real?" The answer is no longer relevant. What now worries voters, journalists and regulators is **any statement coming from a digital entity that can be reproduced infinitely**. When an avatar says "I promise to cut taxes for small businesses," the truth of the statement depends as much on the content as on who — or what — generated that message. **Verifiable provenance**, which the blockchain proof of existence guarantees, becomes the key criterion for authorizing or repudiating a digital replica.

This new frontier sits at the heart of the liar's dividend Beth Noveck describes: "the ability to dismiss true but inconvenient facts as false" [2]. AI avatars amplify this phenomenon; their skill at generating coherent, persuasive speeches makes it easy for humans to delegate their voice to a faceless entity. The moment **the voter cannot distinguish whether the message comes from the real candidate or their avatar**, the election becomes a battlefield of digital authenticity rather than of policies.

This is where the proof of existence (t2p) comes into play. By registering every official statement — text, audio and video — with an immutable cryptographic fingerprint on a **neutral** chain, a public, unalterable witness is created of who spoke and of what was said. This anchor depends on no trusted third party; its total neutrality removes the need for the creator to maintain credentials or for watermarks to be deemed sufficient [10]. As we have seen, initiatives like the *Tech Accord* and Utah's synthetic media legislation recognize that provenance marks are only a minor complement, unable to resolve the trust crisis on their own [5][4].

Yet the solution is still partial. Chapter 81 details the **digital replica rights** that let a person demand the removal or correction of content disseminated without their consent. When an AI avatar repeats an erroneous statement, the victim can invoke that right — but only if they hold unequivocal proof of the original source. Without the **proof of existence**, the chain of responsibility breaks and the debate drifts into endless litigation.

So the question left open is how to integrate this verification layer into electoral processes without deceiving the public with a new form of digital bureaucracy. The answer runs through **turning every electoral message into an immutable record** and, at the same time, establishing legal mechanisms that recognize that record as binding evidence before the authorities. Only then will we be able to say that democracy has found a tool to make the challenge of the political avatar governable, keeping trust in electoral truth alive while preserving digital freedom of expression.

The promise and the peril

On an election day in the Peruvian capital, an observer filmed with her phone a voting table where the ballot in the box and the polling station number were clearly visible. A few minutes later, that image was uploaded to social networks and, without any watermark or label certifying its origin, some political groups began to question it as a "photomontage." The same scene, reproduced with a generative AI, could have been debunked within seconds had the proof of existence registered on the blockchain been able to prove that that instant had been fixed immutably and neutrally. This is t2p's central promise: turning "you can't prove anything" into "this existed and has not been altered," offering an independent anchor that depends neither on a trusted third party nor on credentials that can be revoked.

The impact of this anchoring is immediate. When the truth can be corroborated with an immutable hash, the "liar's dividend" — the power to dismiss inconvenient facts by claiming they are false — loses its force ([2]). Instead of a fruitless race to detect deepfakes, elections become a verifiable proof: the mere existence of the record is enough to deactivate the "maybe it never happened" argument. This not only strengthens voters' trust, but gives the media and the authorities a solid starting point for any subsequent investigation. The blockchain's total neutrality — with no central entity able to manipulate the record — ensures the proof is accepted in courts and in international electoral processes alike.

But this promise comes with an unavoidable danger. First, watermarks and provenance standards, well-intentioned as they are, cannot solve the problem alone ([10]). Without robust media literacy — as the Korean study that reduced the intention to share false information shows ([11]) — the population remains vulnerable to manipulated narratives. Technology can only provide evidence; it cannot replace human judgment or collective trust in solid institutions.

Second, techno-solutionism can deceive with the idea that an immutable anchor guarantees the truth of the content. t2p only certifies that a document or an image existed and has not been altered afterward; it does not assess whether the statements it contains are truthful. Thus an electoral tally could be perfectly anchored yet full of counting errors or human manipulation prior to its capture. The technology must complement verification processes and independent audits, not substitute for them.

Finally, who controls the chain? The critical chapter 112 warns that the infrastructure that guarantees the truth can become a surveillance tool if left without democratic governance. Without clear audit mechanisms and citizen participation, the blockchain itself could be used by authoritarian or corporate entities to record and monitor political behavior, consolidating a new kind of centralized power. This underlines the need for a legal framework that recognizes the record as binding evidence while also limiting abuse.

In short, t2p's neutral anchor is an indispensable tool for restoring electoral trust in an environment where deepfakes constantly erode the truth. Yet its effectiveness depends on solid media literacy, on human processes that verify the content and on democratic governance of the blockchain infrastructure. Only then can we say the tool is necessary — and not merely sufficient.

It's time to act

Electoral authorities are the first guardians of public trust. They must anchor every official act — the vote tally, the preliminary count and the results statement — with a blockchain proof of existence (t2p) and publish its verifiable hash on the electoral portal. This turns any manipulation attempt into a change impossible without a trace, and lets any observer check the document's integrity with a single query to the chain.

The media must adopt the same rigor. Every report including sensitive material — debate videos, recorded statements or leaked documents — must be sealed with a proof of existence and publish its identifier at the end of the article. This gives readers immutable proof that the content has not been altered since it was uploaded, and equips them to verify it without depending on a trusted third party. Journalists must also train their newsrooms in the use of AI-based checking tools that integrate blockchain, as the Voto Libre project in Peru demonstrates [6], where the combination has enabled "secure, transparent, real-time" reporting.

The public, for its part, becomes the last line of defense. Faced with an "explosive" video circulating on the networks during a campaign, the voter should ask for its proof of existence before sharing it. If no verifiable hash exists, the content should be treated as a potential deepfake and reported to the relevant platforms. Moreover, any direct testimony — a recording of one's polling station or a photo of the results board — should go immediately onto a public blockchain, as Taiwan did with the Cofact initiative led by

Audrey Tang [7]. This turns the individual act into a neutral anchor that reinforces the collective fabric of the truth.

Regulators must establish legal frameworks that guarantee the interoperability and verifiable provenance of electoral content, drawing on Utah's example with law SB 131, which demands "tamper-evident" labels for synthetic media [4]. Yet this regulation must not delegate all control to the technology; it must be combined with media literacy and AI training programs, as Beth Noveck recommends: watermarks alone do not solve the problem [10].

In short, the chain of trust is built step by step: authorities anchoring official data, media publishing with an immutable seal, citizens verifying before sharing and regulators imposing standards without abdicating democratic control.

Elections are not taken on faith; they are checked. It's time2act.

---

[1] FCC — *Cloned-voice robocalls and elections (the New Hampshire case, 2024).* https://www.fcc.gov/document/fcc-makes-ai-generated-voice-robocalls-illegal ↩

[2] Beth Simone Noveck — *Reboot: AI and the Race to Save Democracy* (Yale University Press, 2026), ch. 3 ("Ballots and Bots"). ↩

[3] Munich Security Conference — *AI Elections Accord* (2024). https://securityconference.org/en/aielectionsaccord/ ↩

[4] Utah State Legislature — *SB 131, Synthetic Media / disclosure of synthetic media* (2024). https://le.utah.gov/~2024/bills/sbillenr/SB0131.pdf ↩

[5] Microsoft / MSC — *Tech Accord to Combat Deceptive Use of AI in 2024 Elections* (16-02-2024). https://blogs.microsoft.com/on-the-issues/2024/02/16/ai-deepfakes-elections-munich-tech-accord/ ↩

[6] Infobae — *Blockchain y elecciones 2026 en Perú* (06-06-2025). https://www.infobae.com/peru/2025/06/06/blockchain-y-elecciones-2026-como-se-desarrollara-en-peru-el-voto-digital-obligatorio-para-militares-y-ciudadanos-en-el-extranjero/ ; and *Reboot*, ch. 3 (Voto Libre). ↩

[7] Beth Simone Noveck — *Reboot* (2026), ch. 3 (Taiwan, Audrey Tang, Cofact). ↩

[8] The Washington Post — *AI mayoral candidate VIC, Cheyenne* (2024) https://www.washingtonpost.com/technology/2024/08/19/artificial-intelligence-mayor-cheyenne-vic/ ; Wikipedia — *Jason Palmer (politician)* https://en.wikipedia.org/wiki/Jason_Palmer_(politician ; TIME — *Albania's AI minister Diella* (2025) https://time.com/7324934/albania-ai-minister-diella/ ↩

[9] The Markup — *1 in 6 Congresswomen Targeted by AI-Generated Sexually Explicit Deepfakes* (11-12-2024). https://themarkup.org/artificial-intelligence/2024/12/11/1-in-6-congresswomen-targeted-by-ai-generated-sexually-explicit-deepfakes

[10] Beth Simone Noveck — *Reboot* (2026), introduction and ch. 3 ("possibilist, not Pollyanna"; provenance "alone will not solve the problem"). ↩

[11] Beth Simone Noveck — *Reboot* (2026), ch. 3 (the Korean media literacy and deepfakes trial). ↩

[12] IPI / VSquare / HKS Misinformation Review — *the Šimečka–Tódová audio deepfake in Slovakia* (28-09-2023; the 48-hour moratorium; the hole in Meta's audio policy; the debate over the real influence on the result). https://ipi.media/slovakia-deepfake-audio-of-dennik-n-journalist-offers-worrying-example-of-ai-abuse/ · https://vsquare.org/slovak-election-targeted-by-pro-kremlin-deepfake-hoax/ · https://misinforeview.hks.harvard.edu/article/beyond-the-deepfake-hype-ai-democracy-and-the-slovak-case/ ↩

Auditable Public AI: The Record That Can Be Checked

"Models are opinions embedded in mathematics."

— Cathy O'Neil, Weapons of Math Destruction (2016)

It was a cold October afternoon in London when James Whitaker, deputy manager of a regional office of the UK Post Office, received the official notification that would change his life. An email, sent from the corporation's legal department, announced an accusation of fraud against him by a program called Horizon, developed by Fujitsu and deployed in the postal transaction management service. Whitaker, who had devoted more than three decades to postal logistics and enjoyed an impeccable reputation among colleagues and citizens, now stood accused of a crime he had never committed.

The Horizon program, by design, recorded every movement of money across its 600 branches. But a series of software errors created nonexistent "shortfalls": fictitious transactions adding up to millions of pounds. With no independent record that a third party could consult or verify, the software's mere say-so became overwhelming evidence. Whitaker and more than 900 sub-postmasters were unjustly convicted, losing their homes, their reputations and, in some cases, life itself; the case has been linked to at least thirteen suicides [1].

The public inquiry chaired by Sir Wyn Williams, concluded in December 2024, revealed that the absence of an auditable record had allowed the system to be manipulated without any visible trace. This was not just a technical failure, but a structural lack: no one — neither government officials nor the postal agency itself — could reconstruct step by step how those supposed errors had been generated. The result was a catastrophe of trust and accountability that shook the British public sector.

In this context, Noveck's voice in *Reboot* rings clear: "AI must be bought as civic infrastructure — verifiable, auditable and accountable to the people it affects" [1]. The lesson is not to avoid artificial intelligence, but to demand that every machine-assisted decision be traceable and independently checkable. A tamper-evident record, anchored with a proof of existence (t2p), would be the minimum foundation for any future audit; without it, automated decisions remain in the dark, inaccessible to citizen review and vulnerable to impunity.

Thus the scene of Whitaker, an honest man turned victim of an invisible algorithm, illustrates the most insidious danger of delegating public decisions to an AI without a clear accountability framework: irresponsibility. It is not just that the program was flawed; it is that there was no mechanism to check what had really happened and who was responsible. Noveck's proposal — log export for independent review, third-party escrows or restricted evaluator access — thus becomes an urgent necessity for restoring public trust in government technologies [2].

In the sections that follow, we will see how an immutable blockchain-based record can guarantee neutrality and non-repudiation, turning every action of public AI into a verifiable, auditable event, and opening the door to a democracy genuinely accountable before the new forms of automation.

What is at stake: from bias to irresponsibility

Governments already rely on artificial intelligence systems to allocate social benefits, award public contracts and set urban policy priorities. In public debate the focus often falls on algorithmic bias: if an AI learns from biased historical data, it can perpetuate invisible discrimination. But the real wake-up call runs deeper. When decisions are made by a machine and no one can reconstruct the path that led to them, the very essence of accountability disappears. As Noveck says, we must avoid "accountability in name only": a democratic theater where explanations are merely plausible narratives, but not verifiable [9].

This irresponsibility becomes especially dangerous in public procurement, which is the most powerful lever for shaping the economy and public trust. Noveck insists that AI purchasing must demand "log export for independent review" and, when the system is closed, the creation of an escrow or restricted access for external auditors, so that any result can be reproduced and verified [2]. Without these mechanisms, regulations become mere formalities: vendors can store logs on proprietary servers, and administrations can modify or delete data without a trace.

This is where the immutable blockchain-based record takes the leading role. A chained block acts as a neutral witness, impossible to alter for either the administration or the vendor. Every AI action — from selecting a bid to calculating a benefit — is noted with a proof of existence (t2p) that fixes the what and the when, without revealing the how. This minimal information is, on its own, the indispensable foundation for any audit: checking that the event existed, though not necessarily knowing the model's internal logic. Neutrality and non-repudiation are thus guaranteed, while the confidentiality of intellectual property and sensitive data is preserved.

The public algorithmic registers that emerged in Helsinki and Amsterdam in 2020 — and have since been replicated in Canada and Spain — already show how a transparent cataloguing of algorithms can strengthen public trust [4]. Yet without an immutable mechanism, these inventories are vulnerable: they can be edited or deleted, and their value is limited to a declaration of good faith. Integration with blockchain transforms this register into reliable forensic evidence, consultable by independent auditors, courts and interested citizens.

In short, the transition from the bias framing to the accountability framing involves three key steps: (1) requiring log export and auditability as a purchasing condition; (2) establishing escrows or restricted access for external auditors when the systems are closed; and (3) anchoring these records in an immutable infrastructure that guarantees neutrality and non-repudiation. Only then can public AI stop being an instrument of opacity and become verifiable civic infrastructure, as Noveck proposes: "verifiable, auditable and accountable to the people it affects" [1]. This is the base on which we can build a democracy genuinely accountable before growing automation.

The record that can be checked

At the GovLab's laboratory in Washington, the first version of **Policy Synth** was deployed in the summer of 2024 as an open tool for automated decisions in fields as diverse as public tenders and urban policy design. Every action by an AI agent is recorded in a detailed log indicating who acted, what decision was made, when and which data sources fed it; each entry, moreover, comes with an automatic citation of the data used, as Noveck describes: "it keeps a detailed record of its agents' actions, allowing policymakers to audit and explain the decisions; all outputs are referenced and cited" [3].

What turns this record into true **civic infrastructure** is anchoring it with the proof of existence described in chapter 13. When each entry is "timestamped" through an immutable blockchain, the record moves from being a simple digital file to becoming **non-repudiable** evidence: no one — neither the agency that deployed the AI nor the model's vendors — can retroactively modify the recorded information. This eliminates the danger of diffuse responsibility and replaces it with clear, verifiable traceability.

This neutrality is key because, as Noveck notes, the real challenge is not the model's biases but **irresponsibility**: "who takes responsibility when an autonomous AI agent recommends an action that ends in a harmful outcome?" [3]. An immutable record answers this question directly. If a public procurement decision led to a cost overrun or the selection of a corrupt vendor, the log shown by Policy Synth indicates exactly which agent recommended the action, what data backed it and at what moment it was generated. An auditor — or even a court — can reproduce this scenario without depending on one party's good faith or the availability of source code.

As we have seen in Helsinki, the 2020 municipal algorithm register was linked to a proof-of-existence layer to guarantee that decisions on social housing allocation were fully verifiable [4]. In Amsterdam, the register of urban surveillance policies follows the same pattern: every alert generated by a predictive model is timestamped and anchored on the chain, letting citizens request a complete, immutable explanation. These experiences show how the record's **neutrality** — the chain does not judge the content, it only fixes the fact — becomes the fundamental element of a real audit.

In short, anchoring Policy Synth's log with a proof of existence transforms a simple operational history into an indisputable source of truth. Noveck's question about responsibility stops being abstract and becomes a concrete answer: the record indicates **who**, **what**, **when** and **with what data**. This gives public institutions a self-control mechanism that depends not on the vendor's self-sufficiency, but on a public, verifiable, auditable infrastructure — exactly what Noveck defends as "verifiable, auditable and accountable to the people it affects" [1]. Democracy can thus move toward an era where trust rests on unalterable evidence, not just on promises.

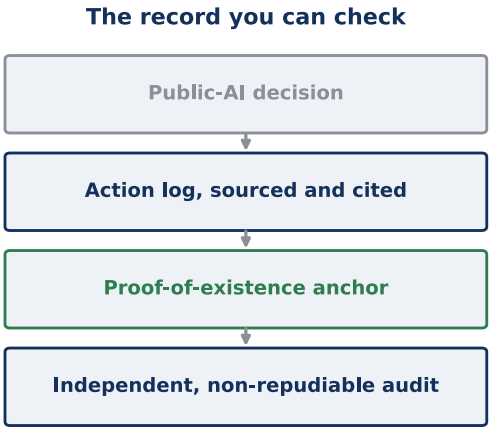


Figure 107.1 — The public AI decision leaves a cited record which, anchored with a proof of existence, becomes an independent, non-repudiable audit.

Registers and provenance: from the catalogue to the proof

In the summer of 2020, Helsinki and Amsterdam published the first municipal algorithmic inventories, documents listing every AI system deployed in public services, from waste management to traffic control. Those catalogues were not mere lists; they included the model's purposes, the training datasets and the human oversight measures applied. The aim was clear: to transform "trust in the city's website" into a trust grounded in verifiable evidence [4]. When these entries were anchored with the proof of existence described in chapter 13, any citizen or external auditor could prove that a record had been created at a given moment and had not been modified since. This removes the option of "purposeful editing" of the logs after a controversial decision and offers a neutral basis for review.

The European efforts inspired similar initiatives in Canada, with the Tracking Automated Government project, and in Spain, where the Algorithmic Observatory has consolidated a national repository covering the AIs used in areas such as public procurement and public safety. These platforms become starting points for independent audits: auditors can download the record, check its time with the proof of existence and compare it with the system's actual decisions. The combination of public inventories and cryptographic anchoring creates a layer of transparency that overcomes the mere "double opacity" Noveck describes, where scientific and commercial knowledge is kept separate and hidden [5].

At the model provenance scale, the AI-BOM (Bill of Materials) concept has gained traction thanks to Noveck's guidelines, which insist on including not only the model's weights but also the training data and the complete infrastructure configuration [8]. DeepSeek R1, released in January 2025 with open weights under an MIT license and a development cost of $5.6 million, has shown it is possible to publish the entire supply chain of an advanced AI without compromising essential intellectual property [7]. When this AI-BOM is registered through the proof of existence, any entity can verify that the declared data is accurate and that nothing was added or removed after training. This provides the minimum foundation on which an external audit can build: both the ingredients and the moment they were combined are known.

The California regulatory framework, with SB 53 and the creation of CalCompute, already obliges large providers to publish standardized safety reports and to maintain a public cloud where executions are immutably logged [6]. This public infrastructure provides the "escrow" Noveck asks for in closed systems: the execution core's logs are accessible to independent auditors, guaranteeing complete traceability from the moment a recommendation is generated to its final deployment.

In short, the so-called "double opacity" is no insurmountable obstacle if a public algorithmic catalogue is combined with the proof of existence and a complete AI-BOM. The register moves from a mere declaration to unalterable evidence that lets both the citizen and the institution verify what existed and when, opening the way to real audits and effective accountability. This is the key for public AI to become civic infrastructure that is verifiable, auditable and, finally, accountable.

The promise and the peril

Imagine a city where every tender, every social policy decision or every police alert generated by an algorithm is registered on an immutable blockchain. An official can open the application, see the transaction with its hash, check that the AI model existed and when, and consult the metadata indicating which weights and data were used. This transforms blind trust into direct verification: "trust the administration" becomes "check it yourself." The proof of existence (t2p) offers, as seen in chapter 13, the minimum foundation for any independent audit; the mere declaration becomes unalterable evidence that citizens and external entities can review.

This capacity to make every AI-assisted action auditable is the central promise of the tamper-evident record. The Policy Synth initiative, developed by Noveck's GovLab, already shows how a detailed log lets policymakers trace and explain an autonomous agent's decisions, with all outputs cited and linked to verifiable sources [3]. Moreover, the California rules (SB 53) oblige "frontier" model providers to publish safety reports and make usage records available to the public, creating an ecosystem where control can be exercised from outside the creating companies [6].

Yet the same infrastructure that guarantees transparency can breed grave dangers if not accompanied by external audit and human oversight. The record's mere existence does not explain the model's provenance or its justice: the problem of "double opacity," where the science is incomprehensible and the commerce stays closed, persists even if every transaction is immutable [5]. Without independent audit mechanisms, the record can become mere "accountability in name only," a democratic theater where explanations are plausible but not true [9].

Another danger appears when trust in the registered information is confused with the decision's legitimacy. The Diella case, Albania's AI minister who promoted public procurement "100% free of corruption" with a human sign-off at every phase, illustrates how an impeccable chain can hide unjust decisions if the model's underlying criterion is flawed [10]. An immutable record cannot correct a policy already resting on a bias or a programming error; it can only prove that the decision was made and when.

Finally, the question of who controls the chain and the nodes is fundamental. Chapter 112 details how the same immutable record that holds the government accountable can, if badly governed, become a tool of mass surveillance over citizens, tracking every interaction with the AI and potentially restricting digital rights. The infrastructure is therefore necessary but not sufficient: data governance, third-party escrow mechanisms and a legal framework must ensure that control of the record stays in neutral, democratic hands [2].

In short, anchoring public decisions in a proof of existence offers the key to real audits and effective accountability. But this promise is reached only if it is complemented with external audit, human oversight and clear governance of the record itself; otherwise, the risk of democratic theater and excessive surveillance can eclipse any expected benefit.

It's time to act

The trust democracy depends on cannot remain an abstract promise; it must be verifiable and, above all, reproducible. That is why every actor in the public sector must take concrete steps to turn public AI into genuinely auditable civic infrastructure.

Public buyers — central administrations and local entities alike — must make mandatory, in their contracts, the inclusion of a record anchored on blockchain and proven by *t2p*. This record must be exportable for independent review and, when the system is closed, a third-party "escrow" or restricted evaluator access must be established, as Noveck recommends: "log export for independent review" and "third-party escrow or restricted evaluator access" are indispensable conditions to guarantee that no decision is lost in an irreproducible black box [2]. The contractual clause must provide penalties if the vendor fails to comply, thus avoiding the Horizon scenario, where erroneous data led to unjust convictions and up to thirteen suicides [1].

Administrations — as the owners of these platforms — must publish algorithmic inventories with every entry anchored to the immutable record. The pioneering examples of Helsinki and Amsterdam, launched in 2020, show it is possible to detail training data, use, human oversight and bias assessment in an accessible public catalogue [4]. These inventories must explicitly name the models' provenance (AI-BOM)

and the data used, as described in chapter 79, so any auditor can trace the origin and the exact moment a decision was generated.

Auditors and courts — as the guardians of accountability — must demand the anchored fingerprint as the minimum basis of review. The proof of existence not only certifies that an action existed; it establishes the exact moment from which the decision process can be reproduced, turning inspection into a feasible task and avoiding the "accountability in name only" where explanations are merely staged without real foundation [9].

Citizens — must have a recognized right to reconstruct and challenge any algorithmic decision that affects them. This right materializes only if the records are public, tamper-proof and accessible for independent review. The legitimacy of a public policy is not self-declared; it is checked against the record showing what steps were followed and when they were executed.

In short, the combination of **neutrality** (an immutable record no one can alter) and **auditability** (the capacity to access, export and review the records) turns public AI into an accountable tool. This promise materializes only if every actor takes on their obligations: contracts with export and escrow clauses for buyers; anchored inventories for administrations; the demand for the fingerprint for auditors; and the right of reconstruction for citizens. Legitimacy is not proclaimed, it is checked. It's time2act.

---

[1] Beth Simone Noveck — *Reboot: AI and the Race to Save Democracy* (Yale University Press, 2026), ch. 8; and *Post Office Horizon Inquiry* (Sir Wyn Williams, 2024). https://en.wikipedia.org/wiki/British_Post_Office_scandal ↩

[2] Beth Simone Noveck — *Reboot* (2026), ch. 8 ("Using procurement power"). ↩

[3] Beth Simone Noveck — *Reboot* (2026), ch. 6 (Policy Synth, runtime guardrails); The GovLab. https://thegovlab.org/ ↩

[4] VentureBeat — *Amsterdam and Helsinki launch algorithm registries* (2020) https://venturebeat.com/ai/amsterdam-and-helsinki-launch-algorithm-registries-to-bring-transparency-to-public-deployments-of-ai ; and *Reboot* (2026), ch. 8 (Canada, Spain). ↩

[5] Beth Simone Noveck — *Reboot* (2026), ch. 8 ("double opacity"; external audits). ↩

[6] Government of California — *Transparency in Frontier Artificial Intelligence Act (SB 53)* (29-09-2025). https://www.gov.ca.gov/2025/09/29/governor-newsom-signs-sb-53-advancing-californias-world-leading-artificial-intelligence-industry/ ↩

[7] DeepSeek — *DeepSeek-R1 release* (20-01-2025). https://api-docs.deepseek.com/news/news250120 ↩

[8] Beth Simone Noveck — *Reboot* (2026), ch. 8 (open source: the five elements; the record of queries and training runs). ↩

[9] Beth Simone Noveck — *Reboot* (2026), introduction and ch. 6 ("accountability in name only"; AI as an object). ↩

[10] TIME — *Albania's AI minister Diella* (2025). https://time.com/7324934/albania-ai-minister-diella/ ↩

Authentic Participation: One Person, One Voice, Without Surveillance

"We do not say that a man who takes no interest in politics is a man who minds his own business; we say that he has no business here at all."

— Pericles, Funeral Oration (Thucydides, II, 40)

It was the summer of 2017 when the Federal Communications Commission (FCC) launched a public consultation on net neutrality. The online form filled with more than **22 million** comments; the noise flooding the platform was so intense that, a few months later, an FCC report revealed that **18 million** of those messages were fake and **9.3 million** carried invented or stolen identities. Most of those submissions came from a single 19-year-old who had programmed a bot to flood the debate, while close to **500,000** comments arrived from Russian addresses linked to a covert campaign financed by the broadband industry [1].

That scenario, as concrete as it is emblematic, showed the paradox that now repeats itself in every digital sphere: when the noise overwhelms the signal and we cannot distinguish a human voice from an algorithm, citizen participation turns into a **theater of voids**. Democratic platforms, which should be stages for authentic expression, end up becoming spaces where "opinion" is just a surface layer of fabricated information. Trust disintegrates; the citizen's voice, with no verifiable proof behind it, loses all value and becomes a murmur no one listens to.

In this context, the need for a **personhood proof** — one person, one voice — becomes imperative. It is not only necessary to guarantee that every vote or endorsement of an initiative comes from a real individual; that guarantee must also avoid depending on a centralized biometric database, which could become an instrument of mass surveillance. The solution must rest on the **proof of existence** and the **on-device proof of life**, as described in chapter 13: a digital signature combining a video selfie with biometric verification performed entirely inside the phone itself, with no external entity accessing or storing the biological data [7].

Unlike systems such as Worldcoin, which scans the iris to "prove you are human" and has been the object of bans and fines in more than nine jurisdictions, including the declaration of unconstitutionality in Kenya and the forced destruction of the collected data in Colombia [4], this decentralized approach guarantees that the personhood proof is **neutral** and **non-centralized**, preventing any company or government from using it as an entry point for surveillance.

So the challenge Noveck poses — scale, action and institutionalization of digital participation — can only be tackled if we overcome the fundamental obstacle of **indistinguishable noise**. Only with a verifiable personhood proof, without depending on a centralized biometric infrastructure, can we transform digital participation from a simulacrum into a real force that decides and does not merely shout. This is the first step toward a democracy in which every citizen has an authentic voice, recognized and respected.

The two failures of participation

In recent years, digital platforms have turned the citizen's vote into a spectacle: massive surveys that gather thousands of voices but never come to transform any political decision. This phenomenon, which Noveck describes as "participation without integration," becomes a theater where democracy acts only as the stage of an empty performance [6]. Open consultations, online petitions and digital referendums are often collected by entities with no mechanism to channel those inputs toward legislative or executive power. The result is an accumulation of noise that dissipates in the crowd and disappears without a trace or real impact.

To this first failure is added another, more insidious one: the escalation of fake identities and bots that drown the debate. An FCC report indicates that, of more than 22 million comments submitted through the net neutrality docket, some 18 million were fake [1]; the New York Attorney General confirms that 9.3 million of those messages came from fictitious identities controlled by a single 19-year-old with automated software and a network of almost 500,000 Russian addresses [1]. As we have seen, this flood of falsehoods distorts the perception of citizen consensus and discourages authentic participation: real users are forced to fight against a tide of deceptive information and, finally, abandon the process.

The danger lies not only in the abundance of noise; it is the loss of trust this practice breeds. A study cited by Noveck indicates that 51% of Europeans surveyed would be willing to replace some politicians with artificial intelligence, a figure that betrays a deep exhaustion of faith in democratic institutions [8]. This statistic is not just a barometer of dissatisfaction; it is a clear signal that the current model of digital participation has lost legitimacy. More noise, more open platforms and more consultations without consequences only worsen the perception that individual voices are irrelevant.

To escape this loop, two basic requirements must be redefined. First, participation must decide: citizen inputs must be connected directly to real decision processes, as the Brazilian Senate's e-Cidadania portal demonstrates, where a "Legislative Idea" that reaches 20,000 endorsements passes to the Human Rights Commission and can become law — a recent example is the "Joca Law," passed in April 2025 after a massive popular movement [2]. Second, that decision must come exclusively from real people. The personhood proof verified by t2p offers a solution: a certificate combining a video selfie and on-device biometric verification with no third party storing the data [7]. Unlike Worldcoin's centralized iris model, banned and fined in nine jurisdictions for violating privacy [4], this proof keeps the identity's integrity at the local scale, avoiding mass control and guaranteeing that every voice is authentic.

In short, the current failure of digital participation stems both from the fact that it does not decide and from the fact that it sinks into a mass of artificial identities. The way out runs through two indispensable pillars: a direct connection to the structures of power and a decentralized verification ensuring that every contribution comes from a real person. Only then can we transform the indistinguishable noise into a democratic force capable of deciding, and not merely of shouting.

Participation that decides

In the Brazilian Senate, the e-Cidadania portal has turned the citizen's voice into a parliamentary obligation whenever a proposal reaches a threshold of twenty thousand endorsements [2]. This is no small show of interest: the case known as the "Joca Law," born of a popular complaint about a dog killed by an airline's error — as we have seen — became law in April 2025 [2]. Today the system records more than a hundred and twenty thousand legislative ideas and tens of millions of endorsements, turning the mass of digital comments into a concrete normative force.

In Taiwan, the vTaiwan model combines open deliberation with binding decisions. The "plural technologies" described by Tang and Weyl connect citizen debates with the institutions, so that proposals which cross a preset threshold of support are officially taken up by the executive [3]. This is not merely a platform where people comment; it is an infrastructure that translates digital consensus into real legislative action.

These two examples illustrate a fundamental rule: technology can only empower democracy when there is a clear institutional framework and a firm commitment to respond. Without that framework, even the most sophisticated tools become mere entertainment stages, where voices carry no real weight. The danger is that, as digital participation expands, the noise generated by fake identities and bots threatens to corrupt this process.

Recent data shows how this vulnerability materializes: **as we have seen**, more than 80% of the comments in that FCC consultation turned out to be fake or manipulated [1]. When an initiative depends on a minimum number of endorsements to advance, the effect of artificial identities translates directly into political decisions. An organized group can create thousands of fictitious accounts, artificially inflate support for a proposal and, without reliable verification, make legislators act on an invented consensus.

This fragility is the gateway to the next step: guaranteeing that every endorsement comes from a real, present person. The proof of existence of **t2p** offers exactly that — a signature bound to a video selfie and a biometric verification performed exclusively on the device, with no external entity accessing the information [7]. Unlike the centralized iris model of **Worldcoin**, banned and fined in various jurisdictions for violating privacy [4], **t2p**'s decentralized mechanism keeps the identity secure and invisible to third parties.

So while Brazil and Taiwan show how technology can channel participation into tangible results, its effectiveness depends on one essential component: a reliable verification preventing bots from replacing authentic voices. Only with this layer of trust can we transform digital noise into a democratic force capable of truly deciding, and not simply of shouting.

One person, one voice, without surveillance

In a São Paulo café, beside a screen showing in real time the digital signatures on the e-Cidadania portal, a young activist named Lúcia watched as the number of endorsements for the "Joca Law" passed 20,000 and, at a stroke, became a legislative proposal. The moment was electrifying: a citizen had managed to turn a small complaint into a national initiative only because her vote, her endorsement, was bound to a verifiable, unique identity. But this success would not be possible without a fundamental layer of trust guaranteeing that behind every signature there is a real, present person, and not a bot or a fake identity.

The current crisis of digital participation shows itself in two dimensions. On one side, many online spaces are merely stages where voices rise with no capacity to influence: a "theater" of comments that never reach political decisions [6]. On the other, the proliferation of fake identities and bots has flooded the platforms with artificial endorsements; as we have seen, the FCC found that almost all the comments received in a 2017 public consultation were fake [1]. This double problem paralyzes digital democracy and feeds citizens' distrust.

One apparent solution has been the Worldcoin project, which purports to prove a user is human by scanning the iris and centralizing that biometric on a blockchain. As we have seen, this model has proven dangerous: between 2024 and 2025 it was banned and fined in more than nine jurisdictions; Kenya declared it unconstitutional and ordered the total destruction of the collected data [4]. Centralizing a trait as sensitive as the iris creates a surveillance nightmare, where one entity — public or private — controls biological information that could be exploited for authoritarian or commercial ends.

The neutral path is offered by the proof of existence and of life developed by t2p. This technology combines a video selfie with biometric verification entirely on the device, with no third party seeing or storing the image or the biometric data [7]. When the user captures their face and voice, a unique cryptographic signature is generated that binds the act (vote, comment, endorsement) to a present person. The proof depends on no centralized database; the trust resides in the integrity of the device itself, guaranteeing that the identity is authentic without revealing it.

Moreover, the "personal APIs" Noveck proposes let the citizen control their own identity and data online [5]. Instead of being a platform's product, the individual becomes the manager of their digital profile, deciding what information to share and with whom. This architecture reinforces the crucial distinction between personhood proof — the fact of being a unique person — and identification — who exactly you are. The former ensures that only real people can participate; the latter remains private and is not needed for civic action.

As we have seen, the proof of existence with on-device liveness offers a direct answer to the problem of fake voices without creating a surveillance hub. It is the same neutrality that has sustained the success of platforms like e-Cidadania or vTaiwan: no third party controls the key to the trust, and every vote becomes an authentic declaration of a single human voice. This architecture is the foundation on which we can transform digital noise into real, democratic decision.

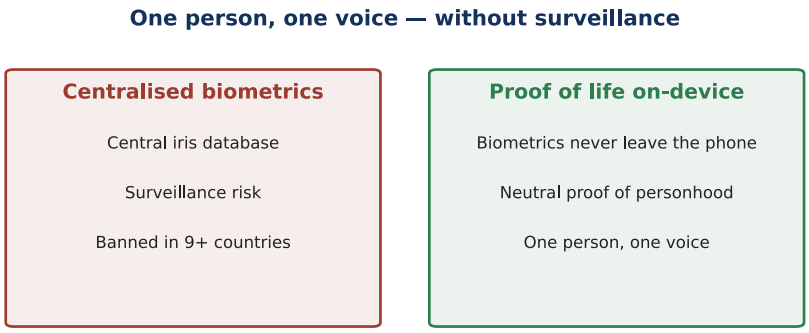


Figure 108.1 — The personhood proof can be performed on the device, with no centralized biometric databases: one person, one voice, without surveillance.

The promise and the peril

Imagine a digital election day where every vote is registered with the security of a physical fingerprint, but with no entity able to access the biometric data that generated it. This is the vision offered by **t2p**'s neutral personhood proof: a digital signature based on a video selfie and a biometric verification entirely **on the device**, binding each act to a real, present person without any third party seeing or keeping that sensitive information [7]. In theory, this mechanism can transform the noise of thousands of fictitious identities into authentic decisions, turning mass participation into a true democratic exercise. The promise is clear: a single human voice per vote, without the blast of bots or the manipulation of fake profiles that has paralyzed the digital initiatives of the past decade.

Yet the road to that full trust is strewn with risks we cannot ignore. First, any personhood proof system runs the danger of becoming a **surveillance architecture** if the biometrics are centralized. The Worldcoin example illustrates this scenario: the proposal to scan the iris to prove a user is human has been banned, fined or ordered to delete data in more than nine jurisdictions, and both Kenya and Colombia have found its model unconstitutional for its dependence on a centralized iris database [4]. This experience shows that a well-intentioned idea can become an instrument of mass control when biometric data becomes an asset managed by external entities. The red line, then, is clear: **never** centralized biometric databases.

Second, the personhood proof is emancipating only if it is integrated into institutional power. Noveck reminds us that digital participation is trapped in three recurring challenges — scale, action and institutionalization; without an effective connection to the decision structures, the result remains a **theater** where citizens speak but do not decide [6]. The successful experiences of Brazil (e-Cidadania) and Taiwan (vTaiwan/Pol.is) show that the key is a neutral architecture where no third party controls the keys of trust, letting every vote become an authentic declaration of a single human voice [2][3]. But this integration requires a regulatory framework guaranteeing that digital signatures are recognized and valued by the legislative institutions.

Finally, the question of **who controls the chain** where these proofs are registered is fundamental. If the underlying infrastructure were managed by private entities or by states with political interests, the emancipating potential could be usurped as an instrument of digital hegemony. Chapter 112 details the limitations of the Worldcoin model and the need for open, democratic governance over any chain that stores personhood proofs.

In conclusion, the neutral personhood proof offers a unique opportunity for authentic citizen participation — one person, one voice — but only if it remains **decentralized**, is integrated into real power and is subject to transparent governance. Without these guarantees, the risk is that the technology, instead of liberating, becomes a new layer of surveillance that domesticates the democracy it purports to reinforce. The final decision rests with us: designing and defending a system where trust is **neutral**, not a currency for control.

It's time to act

Parliaments and city councils must abandon consultations that only collect voices without the power to decide. Mandatory thresholds must be established, like the model of the Brazilian Senate's e-Cidadania portal, where an initiative reaching 20,000 endorsements passes to the Human Rights Commission and becomes a bill [2]. Similarly, Taiwan has integrated the plural technologies of vTaiwan/Pol.is to turn every citizen contribution into a concrete step toward legislative action [3]. The adoption of these frameworks should be accompanied by an on-device personhood proof: a signature with a video selfie and biometric verification performed entirely inside the phone, with no external entity accessing or storing the biometric datum [7]. This mechanism eliminates the bots and fake identities detected in the FCC's 2017 consultation, where 9.3 million comments were fictitious — including a single 19-year-old — and close to half a million came from Russia [1].

Civic designers must recognize that personhood proof is not synonymous with identification. The goal is to guarantee that every vote or proposal comes from a real, present person, but without creating a centralized biometric database like the one Worldcoin tried to impose with iris scanning — a project banned and fined in nine jurisdictions for violating fundamental rights [4]. The lesson is clear: the solutions must be decentralized, minimal and fully controlled by the user's device. Any architecture requiring central biometric repositories must be prohibited by regulation, as the European framework on sensitive data provides.

Citizens must demand participation that decides and tools that let them be participants, not products. The "personal APIs" Noveck proposes offer a control panel where each individual manages their own digital identity without ceding it to third parties [5]; thus the citizen's voice becomes a real force of influence and breaks the vicious circle of "participation without integration" that reduced democracy to a simulacrum [6].

Finally, regulators must protect the decentralized personhood proof and veto any form of central biometric repository. This means legislating against the mass storage of irises or fingerprints off the device and establishing penalties for companies that try to build "surveillance layers" beneath the layer of trust. The regulation must align with the principle of neutrality: technology must be an instrument that reinforces trust, not a currency for control.

In short, a democracy that listens to people — without having to watch them — can transform digital participation into authentic civic action by adopting decision thresholds, implementing on-device personhood proofs and guaranteeing strict regulatory frameworks. It is the moment to apply these concrete measures: time2act.

---

[1] New York State Attorney General — *Report on fake comments in the FCC net neutrality proceeding* (2021). https://ag.ny.gov/press-release/2021/attorney-general-james-issues-report-detailing-millions-fake-comments-revealing ↩

[2] Federal Senate of Brazil — *e-Cidadania portal / the Joca Law* (2024-2025) https://www12.senado.leg.br/ecidadania/ ; and *Reboot* (Noveck, 2026), ch. 7. ↩

[3] Audrey Tang and E. Glen Weyl — *Plurality: The Future of Collaborative Technology and Democracy* (2024); and *Reboot* (2026), ch. 7. ↩

[4] Biometric Update / ICLG — *Worldcoin/World ID: bans and deletion orders (Kenya unconstitutional; Colombia)* (2024-2025). https://www.biometricupdate.com/202510/colombia-orders-world-shut-down-citing-biometrics-compliance-failures ↩

[5] Beth Simone Noveck — *Reboot* (2026), ch. 3 (personal APIs). ↩

[6] Beth Simone Noveck — *Reboot* (2026), ch. 7 (scale, action, institutionalization; "participation without integration"). ↩

[7] time2act — chapter 13 (t2p: the video-selfie signature and on-device proof of life); T2P-Manual (v1.0, 2026). ↩

[8] Beth Simone Noveck — *Reboot* (2026), ch. 3 (the Univ. of Madrid survey: 51% would replace politicians with AI). ↩

Demonstrations: The Symmetrical Proof in the Street

"Injustice anywhere is a threat to justice everywhere."

— Martin Luther King Jr., "Letter from Birmingham Jail" (1963)

Minneapolis, May 25, 2020. At nightfall, in front of a neighborhood store, a seventeen-year-old girl passing by raises her phone and records. Her name is Darnella Frazier, and her video contradicts the initial version of events and becomes central evidence in the trial of officer Derek Chauvin for the murder of George Floyd [2]. On June 11, 2021, the Pulitzer Prize board awards her a special citation for "courageously recording the murder of George Floyd, a video that spurred protests against police brutality around the world, highlighting the crucial role of citizens in journalists' quest for truth and justice" [1]. A teenager with a phone changed history because her video was believed. And this is exactly the question running through this chapter: in the deepfake era, what will happen on the day a video like that can be dismissed with two words — "it's fake"?

This chapter tackles the right to demonstrate and public order: the documentation of the street, in both directions. Evidence in armed conflicts and human rights is in chapter 20; photojournalism, in chapter 17; elections and the "liar's dividend," in chapter 106; civic participation, in chapter 108. The key question: how do you reconstruct a police charge, an altercation or an arrest without the stronger side's version simply winning? This book's answer is symmetrical by construction: evidence anchored at capture by both sides — **t2p**: the SHA-512 fingerprint computed on the device while recording, the RFC 3161/eIDAS timestamp and the anchoring on the public blockchain inside an .asice container [14]. Neither the protester nor the officer can rewrite what was sealed. The proof has no side: that is its democratic strength.

The video against the police report

The Spanish case that condenses it arose from a protest of the "Rodea el Congreso" movement, in Madrid. Three riot officers accused a protester of having kicked and punched them; the initial police report, however, recorded only insults, and the accusations kept growing through the investigation and the trial — until the videos appeared, showing no assault at all. The line from the court reporting says it all: "She is not convicted because there was a video" [3]. The epilogue is uncomfortable: the officers who had accused her ended up acquitted — the court recognized there was no evidence of the assaults they had asserted so roundly, but saw no crime in that [4]. Without the video, a probable conviction for her. With the video, no consequences for anyone.

The procedural terrain is worth understanding, without epic. The police report is not an anticipated verdict: if the officers who drafted it do not ratify it at trial, under publicity and cross-examination, it has no evidentiary force on its own [5]. But in practice the officer's word carries weight, and in the administrative sphere — that of fines — the authority's presumption of veracity leaves the citizen nearly defenseless without images. That day, the video leveled the field. By pure chance: someone was recording.

The lesson of the case is not against anyone; it is against chance. The truth of a square cannot depend on someone happening to have their phone in hand, or on the file surviving, or on the judge believing it when the other side insinuates it has been edited. And the symmetrical lesson holds equally: what protects the honest officer falsely accused of a brutality he did not commit? The same answer: evidence anchored at the moment of capture, as standard, for everyone [14].

Recording the police: a right with a fine

In Spain, the right to document the street has lived through a decade of gray zone. Article 36.23 of the Citizen Security Law — the "gag law," of 2015 — penalizes the use of images of officers in police interventions; the Constitutional Court struck down the term "unauthorized," closing the door on prior censorship, and, according to press reports, police forces no longer apply the infraction in its original form [6]. PSOE and Sumar agreed in 2024 on a reform so that the risk to officers from disseminating images must be "demonstrable" and set out with reasons [6]. The wider context explains the distrust: ten years into the law, by March 2025, more than 400,000 fines had been imposed — under the law as a whole, not

just this article — and the third attempt at reform was still in progress [7]; Amnesty International speaks of "ten years of human rights violations and broken promises" [8].

Recording the police in the exercise of public functions, in good faith and without endangering anyone, is a democratic guarantee: Frazier proved it in Minneapolis [1] and the "Rodea el Congreso" video proved it in Madrid [3]. But whoever records is still exposed to punitive uncertainty, and the material they capture is fragile: a confiscated or broken phone, a deleted file, and the evidence dies with it.

t2p's answer is twofold. First, anchoring at capture: if each segment's fingerprint is anchored while recording [14], not even the destruction of the device erases the material's dated existence — what was recorded stands attested, even if a copy has to be recovered later. Second, the protection of the one who records: the fingerprints-only design publishes nothing, identifies no one and disseminates no image; it fixes that this existed, and that is enough. Documenting stops being a risky act with fragile proof and becomes a discreet act with indestructible proof.

The body camera: the gaze from the other side

The symmetry begins with the police forces, and the rollout is already a fact. The Mossos d'Esquadra have extended Axon Body 3 body cameras to 98% of the territory, with 800 units and 300 more planned; 5,000 officers have received the necessary training and authorization hours have been enabled for 7,000 police, with 1,500 more on the way; the usage regulation is being updated, and the unions demand its generalization [9][10]. The National Police is gradually extending its "Personal Recording Device" (the AB3 model) among public safety officers, and it is expected to reach the Guardia Civil too; the declared objective is double: to protect the public and to protect officers against false accusations [11].

But the rollout does not solve the critical point: custody. Who keeps the footage? Who decides which fragment gets published, which one "gets lost," which one accompanies the press release? A body camera without verifiable integrity is, in essence, the memory of one side only: it can document the assault on the officer, and it can vanish when it documents the excess. The trust it generates depends, entirely, on trust in its custodian.

This is where this book's proposal completes what the cameras have begun: a continuous chain of fingerprints, anchored at capture [14]. Every video segment is sealed at birth, so that any cut, gap or later edit becomes mathematically detectable. The honest officer wins first: no one will be able to say his exculpatory video is a fabrication. And the citizen wins at the same time: no selective deletion will pass unnoticed, and the commander who publishes a fragment will have to accept that the context exists, is dated and is judicially demandable. The body camera with anchoring stops being the memory of the stronger side and becomes everyone's memory.

"This is AI": the liar's dividend reaches the street

In August 2024, Donald Trump falsely claimed that the photograph of the crowd waiting for Kamala Harris at the Detroit airport was "AI-generated." The photo had been taken by a campaign staff member and was not manipulated; there were some 15,000 people, and journalists from the Associated Press and other outlets, present on the same tarmac, had documented the crowd from every angle; fact-checkers debunked it in a chain [12][13]. The lesson is ice-cold: if a crowd photographed by dozens of professionals can be denied with two words, what cannot be denied about an inconvenient ten-second video recorded by an anonymous protester — or by a body camera?

The "liar's dividend" — the concept we introduced in chapter 106: the ability to dismiss authentic but inconvenient evidence as fake — cuts in two directions when it reaches the street. In one, power can deny the real video of a police excess: "a fabrication," "AI," "out of context." In the other, a violent group can deny the real video of its rampage — or, worse, fabricate a fake one to criminalize a peaceful protest, with the same synthesis tools we have already seen applied to voices and images in earlier chapters.

In both scenarios, the structural way out is the same: for the authentic material to carry an anchored date and fingerprint from the first second [14], and for material without an anchor to be worth less every day before a court and before public opinion. The truth of the street cannot be defended by shouting; it is defended with mathematics.

The proof of existence applied

A full square, and two clocks recording: a protester's phone and an officer's body camera. Neither of the two accounts will be rewritable.

1. The protester who documents

A protester records a police intervention. While she records, the device computes the SHA-512 fingerprint by segments and anchors it with the RFC 3161/eIDAS timestamp [14]. She publishes nothing: the anchoring reveals neither content nor identity. If the phone is lost, broken or confiscated, the existence and date of the material stand attested. If she is accused of having assaulted someone — as in "Rodea el Congreso" [3][4] — the complete video surfaces with an authenticity no generic challenge can knock down. If someone shouts "this is AI" [12], the fingerprint predating the denial refutes it. And if she committed the assault, the same seal will serve the prosecution: the proof has no side.

2. The officer with the anchored camera

The officer's body camera [9][11] generates a continuous chain of fingerprints: every segment, sealed at capture [14]. Against a false brutality complaint, his video is unassailable. Against a suspicion of selective deletion, the chain shows whether there are gaps, and auditors — internal, judicial or parliamentary — can inspect the whole sequence. The police force that publishes a fragment does so knowing the context is dated and demandable. Transparency stops depending on the goodwill of the moment.

3. The observer and the journalist

The human rights observers and photojournalists covering the demonstration (chapters 17 and 20) work with equipment that seals at capture: their material becomes an evidentiary package with a chain of custody from the street to the courtroom [14]. And they can do it without exposing their sources: fingerprint first; publication, only if needed and when it is safe — the same verifiable evidence model the book showed for armed conflicts in chapter 20.

4. The symmetrical reconstruction

The court receives the two anchored streams — citizens' phones and body cameras — and, instead of word against word, obtains a merged, verifiable timeline: what happened first, who responded with what, with what proportionality. The police report, which on its own needs ratification [5], arrives accompanied by material whose date and integrity admit no argument. Investigations shorten; judgments rest on facts.

The limits, stated with the same clarity: the seal grants no right to record everything in every circumstance — the reform of the citizen security law is still pending [6][7], and third parties' privacy counts; it does not weigh the proportionality of a charge or the legality of an arrest, which belong to the judge; it replaces neither witnesses nor experts; and an authentic clip can still mislead if the context is cropped. Anchoring attests existence and integrity; it does not attest completeness.

The promise and the peril

The promise: that never again should the truth of a square depend on the chance that someone was recording [1][3], or on the unilateral custody of some footage [9], or on credulity before the "this is AI" [12]. Symmetrical, anchored, neutral evidence: the evidentiary infrastructure of a democracy that takes to the streets. The perils, however, are serious, and they must be faced.

The first is the **surveillance inversion**. The tools for documenting power cannot end up identifying those who protest. Anchoring a video must imply no register of who was in the square: the fingerprints-only design reveals neither content nor identity, and facial recognition on demonstration material is a line a democracy must not cross (see chapter 108: participation without surveillance). No central registry, no deanonymization.

The second is **the decontextualized authentic**. The seal attests that a clip existed; it does not attest what happened off-frame. A truthful fragment cropped with malice is the hardest disinformation to fight. The mitigations: continuous chains instead of loose clips, sequence metadata and a culture of context — the seal can never be a total argument from authority.

The third is **the chill over the street**. If demonstrating means knowing yourself recorded from everywhere and forever, fear can empty the squares — the same dampening effect the gag law's sanctions have cast over photographers and activists for a decade [7][8]. That is why the revocability and expiry of personal fingerprints, and always-voluntary publication, are not technical extras: they are democratic conditions of the system.

And the fourth is **the technological alibi**. No camera replaces the pending legal reform [6][7], or police training, or independent accountability mechanisms. The seal makes the truth recoverable; justice, the institutions must deliver.

It's time to act

The square belongs to everyone, and so must its truth.

For protesters and citizens: record in good faith when needed, seal at capture [14] and do not publish third parties' faces without need; know the current legal framework [6] and remember that the fingerprint protects you even if you never publish anything.

For police forces: deploy the body cameras with an anchored chain of fingerprints [9][11][14] and publish the custody and retention policies. The honest officer is the first beneficiary of evidence no one can call into question.

For judges and prosecutors: demand anchored integrity from both sides' material; a verifiable timeline is worth more than a hundred ratifications of a police report [5].

For legislators: complete the reform of Article 36.23 with an explicit right to document in good faith [6] [7]; grant a presumption of authenticity to evidence sealed at capture; and ban mass facial recognition on protest material.

For the media and human rights observers: build the capture seal into your equipment as standard (chapters 17 and 20), verify the anchoring before publishing third-party material and treat unfingerprinted material with the caution it deserves [12][13].

Darnella Frazier was seventeen and had a phone, and the world believed her [1][2]. The next Darnella may not be believed — unless her video carries a date no powerful hand can erase. So that the square keeps telling the truth, time2act.

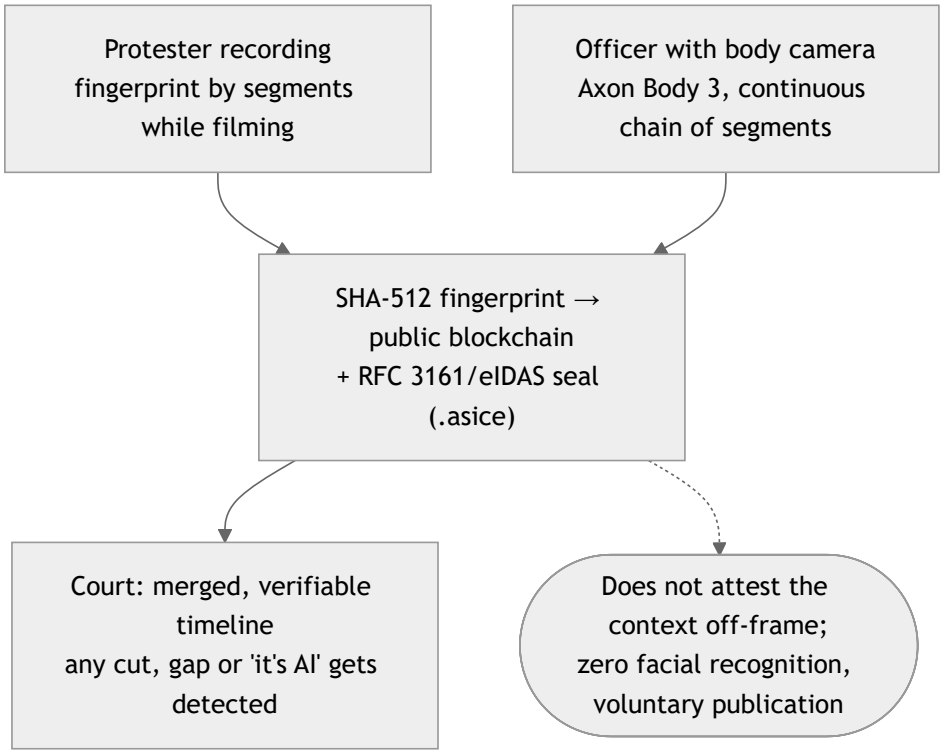


Figure 109.1 — The symmetrical proof in the street: the protester's video and the officer's body camera are sealed at capture, by segments (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, anchoring on the public blockchain, .asice), publishing nothing and identifying no one. In court, the two chains merge into a verifiable timeline; any cut or deletion is detected. Limits: the seal does not attest the context off-frame, does not weigh proportionality and demands revocability, voluntary publication and zero facial recognition. Sources: Pulitzer/Frazier; elDiario.es ("Rodea el Congreso"); SME/USPAC (Axon Body 3); Telemundo/Univision (Detroit).

---

[1] NPR — *Pulitzer Board Recognizes Darnella Frazier For Recording George Floyd's Murder* (June 11, 2021; the special citation; the jury's verbatim wording). https://www.npr.org/2021/06/11/1005601724/darnella-frazier-teen-who-filmed-george-floyds-murder-wins-pulitzer-prize-citati ↩

[2] Wikipedia — *Darnella Frazier* (17 when she recorded; the video as evidence in the Derek Chauvin trial). https://en.wikipedia.org/wiki/Darnella_Frazier ↩

[3] elDiario.es — *Juicio a tres policías por acusar en falso a una manifestante de agresión: «Ella no está condenada porque había un vídeo»* (the "Rodea el Congreso" case; the initial report recorded only insults). https://www.eldiario.es/politica/juicio-tres-policias-acusar-falso-manifestante-si-no-llegamos-encontrar-video-hubieran-condenado_1_10922988.html ↩

[4] elDiario.es — *Absueltos los antidisturbios acusados de mentir en un juicio contra una manifestante del 'Rodea el Congreso'* (the acquittal; the court sees no crime despite the absence of evidence of the assault). https://www.eldiario.es/politica/absueltos-antidisturbios-acusados-mentir-juicio-manifestante-rodea-congreso_1_10979336.html ↩

[5] Togas.biz — *Si el atestado policial no ha sido ratificado en el plenario... las manifestaciones en él contenidas no tienen eficacia probatoria* (the doctrine on the evidentiary force of the police report). https://togas.biz/articulos/articulo-profesionales--si-el-atestado-policial-no-ha-sido-ratificado-en-el-acto-del-plenario-por-los-agentes-de-la-autoridad-que-lo-confeccionaron-someti-ndose-al-principio-de-publicidad-y-contradicci-n-las-manifestaciones-en-l-contenidas-no-tiene-la-eficacia-probatoria/ ↩

[6] Infobae — *El artículo de la 'ley mordaza' que derogarán PSOE y Sumar fue declarado inconstitucional y los policías ya no lo aplican* (July 17, 2024; Art. 36.23; the Constitutional Court strikes down "unauthorized"; the agreed reform: a "demonstrable," reasoned risk). https://www.infobae.com/espana/2024/07/17/el-articulo-de-la-ley-mordaza-que-derogaran-psoe-y-sumar-fue-declarado-inconstitucional-y-los-policias-ya-no-lo-aplican/ ↩

[7] Infobae — *Más de 400.000 multas y un tercer intento de reforma: la ley mordaza cumple 10 años* (March 31, 2025; fines under the law as a whole). https://www.infobae.com/espana/2025/03/31/mas-de-400000-multas-y-un-tercer-intento-de-reforma-la-ley-mordaza-cumple-10-anos-en-medio-de-las-criticas-sobre-su-impacto-en-derechos-y-libertades/ ↩

[8] Amnesty International Spain — *Ley Mordaza: diez años de vulneraciones de derechos humanos y promesas incumplidas* (2025). https://www.es.amnesty.org/en-que-estamos/noticias/noticia/articulo/ley-mordaza-diez-anos-de-vulneraciones-de-derechos-humanos-y-promesas-incumplidas/ ↩

[9] SME — Sindicat de Mossos d'Esquadra — *Consell de la Policia de 6 de febrer del 2025* (Axon Body 3 across 98% of the territory with 800 cameras; 5,000 officers trained; authorization for 7,000 and 1,500 more; the regulation being updated). https://sme-mossos.cat/consell-de-la-policia-de-6-de-febrer-del-2025/ ↩

[10] USPAC — *Resum reunió subcomissió de materials, 27 de novembre de 2024* (300 more cameras planned; requests for generalization). https://uspac.cat/resums-reunions/resum-reunio-subcomissio-de-materials-27-novembre-2024/ ↩

[11] Cámaras de Vigilancia IP — *Cámaras unipersonales de la Policía: funcionamiento y normativa* (the DGU, AB3 model; gradual extension; planned for the Guardia Civil; the aim: protecting the public and the officers). https://camarasdevigilanciaip.es/camaras-unipersonales-policia/ ↩

[12] Telemundo — *Trump afirma falsamente que una foto de la multitud en un mitin de Harris en Detroit fue creada con inteligencia artificial* (August 2024; a photo by a staff member, unmanipulated; AP journalists on the tarmac; 15,000 attendees). https://www.telemundo.com/noticias/noticias-telemundo/elecciones-estados-unidos-2024/trump-afirma-falsamente-que-una-foto-de-la-multitud-en-un-mitin-de-har-rcna166423 ↩

[13] Univision — *Esas fotografías de eventos de campaña de Kamala Harris no fueron creadas con IA, como afirman Trump y en redes* (the fact-check). https://www.univision.com/noticias/fotografias-eventos-campana-kamala-harris-no-ia-como-afirman-trump-y-redes ↩

[14] t2p — product manual (SHA-512 fingerprint on the device, by segments while recording; RFC 3161/eIDAS timestamp; anchoring on the public blockchain; .asice container; a continuous chain of fingerprints for body cameras; optional publication). See chapters 13 and 15. ↩

Digging In: The Cost of Doing Nothing

"I hate the indifferent. I believe that living means taking sides. [...] Indifference is the dead weight of history."

— Antonio Gramsci, "Indifferenti," La città futura (1917)

In 2050, on the first afternoon of a July that seems never to end, the sun rises like an incandescent disk over an uninhabited plain south of the Sahel. The dunes, now coated in layers of fine dust that reflect the light in metallic gray tones, are the only visible trace of a civilization that decided to dig in its heels before climate change. On the border between what remains of Mali and what was once a coastal city, a group of refugees — men and women with their most precious belongings wrapped in worn rags — gathers around a sand well that still holds water. It is not just water; it is the last resource that can save them from certain death by dehydration. In this world, the insurers pulled out long ago from regions they once considered "moderate risk," and coverage against climate disasters has become an unattainable luxury for most.

This scene is not a prediction: it is one of the futures we can still choose or avoid. And what makes it chilling is that in 2026 we already knew enough to prevent it. The IPCC reports showed that keeping global warming below 2 °C would require a decisive energy transition and investments of 1-2% of world GDP [3]. But the previous decade's policies failed to apply these measures, and let emissions continue on their upward path. In that future, warming has passed 3 °C above the preindustrial level, just as the projections for the policies then in force warned [1]. At this level, the economic models indicate that world GDP per capita could fall more than 50% by the year 2100 if we reach 2 °C above the current level [1]; and at 3 °C, cumulative output shrinks between 15% and 34% relative to what was expected [1].

The effects are felt in every corner of the planet. Heat waves stretch on for weeks, strain health systems and send their spending soaring once the illnesses linked to prolonged heat exposure and malnutrition are added. The net cost of inaction comes to represent between 11% and 27% of the world's cumulative output — three times the global health spending projected through 2100 — a burden no economy can bear without consequences [2]. And the climate system's tipping points, once crossed, can cost up to €2.4 trillion a year [2]. Mass migrations have become an everyday reality; entire coastal cities lie flooded and abandoned.

So that 2050 is not just a further year along the timeline; it is the reflection of a collective decision that preferred the comfort of inaction to the energy of transformation. The price of procrastination is measured in lives lost, economies disintegrated and a global trust eroded — a double danger, climatic and of credibility, that feeds on itself. The question left hanging, echoing among the dunes and the abandoned cities, is simple but profound: how much future are we willing to sacrifice in order to stand still?

Where we are heading

The numbers do not lie: current policies are dragging us toward a scenario of **more than 3 °C of global warming**, and this is already a documented reality [1]. The distance between the promises made at the COP summits and the reality on the ground, described in chapter 6, is not just a matter of diplomacy; it is paid in degrees. Every tenth lost represents lives fading away, economic sectors collapsing and a global trust that shatters like glass under pressure.

As we have seen, the projections indicate a fall in GDP per capita of more than 50% and losses of between 15% and 34% of world output through 2100, with costs that could reach up to **€2.4 trillion a year** [2]; these figures reflect an economy crumbling before ever more intense climate phenomena — prolonged droughts, devastating floods and forest fires that raze key infrastructure.

By contrast, the investment needed to hold warming below 2 °C is surprisingly modest: **between 1% and 2% of world GDP** [3]. This figure represents only a fraction of what remaining in our current passivity would cost, without even counting the lost opportunities of sustainable growth built on clean technologies. To assume that acting entails prohibitive costs is therefore an economic and moral error; the true cost is abstention.

Moreover, the double danger — climatic and of trust — feeds back on itself. Every additional degree worsens environmental conditions and fuels distrust in institutions and democratic processes, as the public perceives an incapacity to act in the face of the crisis. This erosion of credibility reduces the

collective capacity to mobilize the resources and political support needed for ambitious measures, generating a vicious circle ever harder to escape.

As we have seen, current policies lead us straight toward a temperature above 3 °C, translating into massive economic losses and rising health spending. The investment to halt this course is comparatively small next to the cost of doing nothing; the task now is to break this inertia, recognize that every degree lost is a warning, and act immediately, for only thus can we preserve both the planet and the trust in our institutions.

The price of doing nothing

Imagine a scale that measures the world's economic well-being like a thermometer: every additional degree of global warming cuts GDP per capita, and the difference is brutal. As we have seen, a rise of 2 °C could reduce this indicator by more than 50% by the end of the century, while passing 3 °C would cause a fall of between 15% and 34% of cumulative output through 2100 [1]. This decline is not just an abstract figure; it means that much of the planet's people would see their capacity to buy, save and invest cut in half, translating into fewer resources for education, health and innovation.

But the direct economic cost is not the only weight on this scale. The net cost of inaction — the difference between the additional expenses generated by climate change and the benefits of an adaptation or mitigation policy — sits between 11% and 27% of cumulative output [2]; as we have seen, this range equals three times the global health spending projected through 2100, revealing how the consequences of climate change will reach every sphere of human life: from the rise in heat-related illnesses to the appearance of emerging epidemics through the migration of vectors.

Moreover, there are tipping points that threaten to break the system at critical moments. Estimates indicate that these structural shifts can cost up to €2.4 trillion a year [2]. These figures are not merely additional costs; they represent the urgent need to invest in resilient infrastructure, clean technologies and governance systems capable of managing the rising risks. Every year that passes without a coordinated response accumulates this hidden cost and turns it into an ever heavier burden for future generations.

In short, inaction is not a neutral state: it is an economic and social bet that costs far more than any preventive investment could ask. The figures show clearly that the price of standing still significantly exceeds that of acting now, both in lost material well-being and in the deterioration of institutional trust. The key to avoiding this desolate scenario is recognizing that every degree lost raises the costs, and cutting emissions immediately to hold warming below 2 °C — a task that takes only 1-2% of world GDP [3]. It is an affordable commitment compared with the colossal bill of inaction.

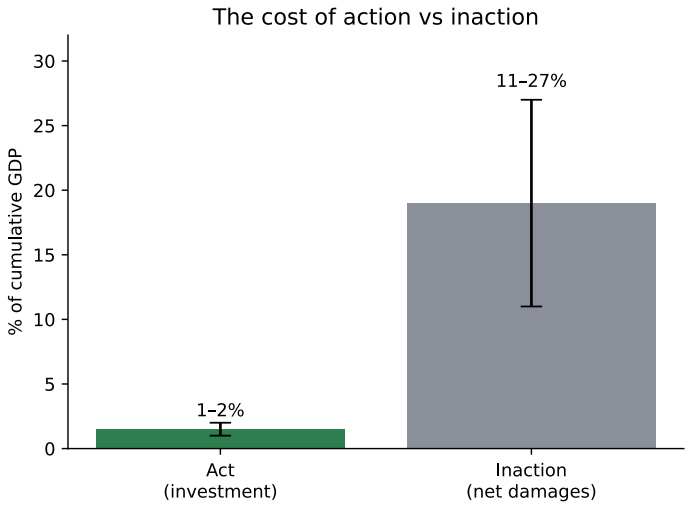


Figure 110.1 — The higher the temperature climbs, the further the economy falls: the cost of not acting soars with every tenth of a degree.

The double systemic risk

At the heart of a crisis that forks between truth and climate, the two entangle in a feedback spiral that amplifies every vulnerability. The uncertainty fed by sophisticated deepfakes, greenwashing reports and manipulated narratives erodes institutional trust and prevents effective governance of the energy transition; without reliable data, climate policies become castles in the air, vulnerable to economic pressures and to disinformation that dilutes the scientific consensus. Every delay increases the climate costs.

Conversely, a warmer scenario generates an even deeper crisis of truth: extreme phenomena — floods, forest fires and prolonged droughts — produce chaotic flows of information that often translate into sensationalist news or falsehoods crafted by particular interests. When communities suffer real damage, public perception fragments; some seek scientific explanations and others adopt conspiracy narratives that minimize human responsibility, hindering governments' application of adaptation or mitigation measures.

The data shows clearly how this double crisis multiplies the global risk; if current policies stay on their trajectory, global warming will pass 3 °C and GDP per capita will fall drastically — as we have seen. The net cost of inaction equals a substantial share of global cumulative output, exceeding three times the health spending projected to the end of the century, and the tipping points could reach up to €2.4 trillion a year if urgent measures are not taken [1][2].

By contrast, containing warming below 2 °C requires only the most modest investment — between 1 and 2% of world GDP [3] — an affordable figure compared with the colossal bill of inaction. This disparity shows that the true danger lies not so much in the technologies or the initial costs, but in the outbreak of a crisis of trust that paralyzes the collective response; reinforcing reliable institutions and transparent verification systems is as urgent as any climate policy.

So the double systemic risk — climate and truth — is not two independent phenomena, but two faces of the same coin. The inability to trust the data prevents governing the transition; the lack of effective governance worsens the climate, which in turn generates more uncertainty and disinformation. Breaking this vicious circle is time2act's central goal: showing that, with artificial intelligence and blockchain to guarantee the integrity of information, we can restore trust and drive decisive climate action before the costs become unbearable.

The cost of the lie, face to face

The cost of not acting is not measured in degrees alone. There is a second bill, quieter, already being paid: that of collapsing trust. And, like the climate one, it grows with every year of delay. Let us walk it across the five fronts this book has examined.

In **finance**, a single forged video call cost the engineering firm Arup some $25.6 million in 2024 [4]; and the consultancy Deloitte projects that AI-driven fraud could reach $40 billion a year in the United States by 2027 [5]. Every scam erodes a little more of the trust that greases commerce.

In **healthcare**, a manipulated medical image — a tumor injected or erased — does not cost money: it costs accurate diagnoses, timely treatments and, ultimately, lives. In **education**, fake degrees and unverifiable authorship depreciate the merit of an entire generation. And in **democracy and defense**, electoral deepfakes and war disinformation shatter the shared base of facts without which no society can agree to act — not even on the climate.

And these five fronts are only the beginning. Throughout the book, the same bill of distrust has kept appearing in every corner of collective life: in the **climate**, where inflated carbon credits and greenwashing empty every saved tonne of meaning; in the **real economy**, from insurance and the supply chain to real estate, asset tokenization, intellectual property and public tenders; and, in this last stretch, in the very fabric of society: the **identity** of those with no document to prove they exist, the **electoral integrity** of tally sheets and counts, the **humanitarian aid** that must reach the real person, **wills**, **international trade**, the **auditing** of accounts, **advertising fraud**, **cybersecurity logs**, **artistic heritage** and the protection of **whistleblowers** who stake their freedom. And further still, on the new frontiers: **science**, where a fabricated datum corrupts knowledge; the **body**, from pensions and doping to the genetic datum that also belongs to your family; the **machines**, which decide crashes and keep their own black box; the specialized **markets** of music, hardware and raw materials; and **justice itself**, where an authentic recording can be denied as fake. On every front, the mechanism of the harm is the same —

evidence that can be fabricated or denied — and so is the defense: a proof of existence no one can alter with impunity.

Because here the two costs converge. A society that cannot agree on what is real will not agree on what must be done about 3 °C. The "liar's dividend" — the ability to discard any inconvenient proof by branding it fake — paralyzes climate action as much as direct disinformation does. The crisis of truth and the crisis of the climate are not two separate emergencies; they are the same emergency seen from two angles, and they feed each other [6].

The conclusion is once again the same, and that is why it is so powerful: the cheapest insurance against both bills is verifiable truth. Anchoring the data — climatic, financial, medical, educational, evidentiary — so that no one can alter it with impunity is not a technical expense; it is the precondition for any other investment to make sense. Without a reliable base of facts, even the 1-2% of GDP needed to curb warming would be spent amid noise and suspicion. With it, every euro counts. The cost of the lie, like that of carbon, only rises while we fail to act.

The rational bet

Imagine a boardroom full of senior officials watching a graph of rising lines: every additional degree of warming accumulates like an invisible burden on the world's economies. At that moment, a voice breaks the silence and, as we have seen, points out that just a small fraction of world GDP invested in the transition could hold warming below 2 °C. This is the promise of science-backed climate policies, and the cost of not acting is a time bomb ticking louder by the instant.

Current projections show that insufficient policies are taking us toward a warming scenario above 3 °C before the century's end [1]. In this context, the economic impact is devastating: as we have seen, if the planet reaches a temperature 2 °C above the current level, gross domestic product per capita could fall more than 50% by 2100, and at 3 °C the world's cumulative output would shrink between 15% and 34% [1]. In parallel, the net cost of inaction is estimated at between 11% and 27% of cumulative output — equal to three times global health spending through 2100 — with tipping points that could reach €2.4 trillion a year [2].

By contrast, devoting between 1% and 2% of GDP to financing the transition to renewable energy sources, improving energy efficiency and deploying systems that guarantee the transparency of climate data represents a comparatively small expense [3]. This investment is not an additional cost, but an economic insurance without precedent: the return is measured in reduced risks and preserved long-term growth. Every euro invested now avoids decades of massive losses, from floods that devastate infrastructure to health crises fed by rising temperatures.

Moreover, investment in artificial intelligence and blockchain offers a double layer of protection. On one side, AI can optimize energy generation and distribution processes, cut unnecessary consumption and increase the system's resilience; on the other, the blockchain ensures that climate data is verifiable and hard to alter, countering the wave of disinformation and deepfakes that undermine public trust. This combination aspires to create an ecosystem where reliable information feeds sound political decisions, breaking the vicious circle between climate uncertainty and loss of credibility.

Ultimately, the rational bet consists in recognizing that the cost of acting is a tiny fraction of world output, while the price of inaction stretches across decades and touches every sector of society. The decision is not between paying and not paying; it is between investing in a sustainable future or letting the economic and social consequences overwhelm us. As citizens, businesspeople and leaders, we have the opportunity — and the responsibility — to contribute this 1-2% of world GDP, and thereby secure a more robust economy and a habitable planet for future generations. The question that remains is: are we willing to make this small bet to guarantee our collective well-being?

The promise and the peril

At the edge of the climate limit, the tools that a few years ago seemed science-fiction fantasies are now tangible, economically viable resources. Renewable energy sources — solar photovoltaic, offshore wind and green hydrogen — have cut their costs spectacularly and, in much of the market, are already cheaper than fossil ones. The transition thus becomes both an environmental necessity and an opportunity to reduce business operating costs.

This new energy abundance is amplified by artificial intelligence, capable of optimizing the balance between supply and demand in real time, minimizing losses and maximizing the use of intermittent resources. Ever more precise weather prediction algorithms are integrated into smart grids, reducing the need for costly reserves and increasing the reliability of supply. In parallel, the blockchain offers transparency and traceability that counter the opacity of traditional energy markets; smart contracts can automate the payment for green energy and ensure that every kilowatt-hour really comes from a clean source, minimizing the risk of greenwashing.

As we have seen, investing just between one and two percent of world GDP would hold global warming below 2 °C, the threshold set by the Paris agreement [3]. It is a modest expense compared with the reduction in losses from climate disasters and with the productivity gains it brings.

Yet the danger presents itself as an invisible but potent force: institutional inertia and cultural resistance to change. Disinformation — fed by social networks and economic interests — distorts the real data and makes the public perceive the solutions as costly or risky when, in reality, they are cheaper than the cost of inaction; **as we have seen** [1].

Moreover, the hidden costs of inaction are already making themselves felt: the estimated net cost of not acting rises to 11%-27% of cumulative global output through 2100, equals three times the world's health spending for that period and could reach €2.4 trillion a year [2]. This impact does not only hit the economy; it erodes social and institutional trust and generates a vicious circle where uncertainty feeds more inaction.

As we have seen, we hold a toolbox of technologies already proven efficient and economical. The challenge is not finding solutions, but overcoming the psychological and political barriers that push us down the most expensive path: inaction. If we choose the latter, the accumulated cost will far exceed that of any responsible investment today. The decision becomes, therefore, a question of vision: do we prefer to invest a small percentage of our gross domestic product to guarantee a sustainable, reliable future, or will we let inertia and disinformation carry us toward a far higher and irreversible cost?

The two curves

Imagine a control room projecting two simultaneous graphs: on one side, the apparently rising curve of world GDP if current policies continue as they are; on the other, the abrupt descent of gross domestic product when warming passes 3 °C. The first shows apparent growth, but over the long run reveals a fall in cumulative output of between 15% and 34% if we reach this critical threshold [1]; the second indicates that the net cost of inaction could absorb between 11% and 27% of global output, the equivalent of three times the world's health spending through 2100, with tipping points reaching €2.4 trillion a year [2] — as we have seen.

By contrast, the investment needed to hold warming below 2 °C sits between 1% and 2% of world GDP [3]. This figure, compared with the shares of output that will be lost if we do nothing, is little more than a budget adjustment. It is, therefore, a decision on a human scale: devoting a small part of our gross domestic product to the energy transition and to the digital infrastructures that guarantee transparency and traceability.

The plan must be as clear as the map of a sea route. First, set intermediate milestones with short deadlines and fund pilot projects that combine artificial intelligence to optimize energy use and blockchain to certify the clean provenance of resources. Second, create accountability mechanisms based on smart contracts: if a company exceeds the established targets, it receives incentives; if it fails, verifiable consequences follow. Third, mobilize private capital through green bonds and impact funds auditable in real time thanks to the blockchain.

This roadmap is not a mere list of measures — it is a collective contract among governments, companies and citizens. Every action must be measurable, verifiable and, above all, adjustable if the results do not match expectations. Trust is rebuilt when the public can see, in real time, how the money is invested and what tangible impact it generates. Thus the double threat — climatic and of trust — is defused simultaneously: an effective reduction of emissions lowers the environmental pressures that feed social uncertainty, and a transparent infrastructure reinforces institutional credibility.

Ultimately, the challenge is as political as it is technological. Leaders must adopt a long-term vision and recognize that the cost of prolonged inaction far exceeds the immediate expenses of a transition plan. The

moment is now; the route is drawn, and all that remains is to decide whether we advance with firm steps or stay trapped in stagnation, letting history record a lost opportunity and an impoverished future.

---

[1] Oxford Economics, *The global economic costs of climate inaction*. https://www.oxfordeconomics.com/resource/the-global-economic-costs-of-climate-inaction/ ↩

[2] Climate Policy Initiative, *The Cost of Inaction*. https://www.climatepolicyinitiative.org/the-cost-of-inaction/ ↩

[3] (a synthesis of climate-economic modeling) Earth.org. https://earth.org/have-economists-been-underestimating-the-cost-of-climate-change/ ↩

[4] World Economic Forum — *Cybercrime: lessons from a $25m deepfake attack* (the Arup case). https://www.weforum.org/stories/2025/02/deepfake-ai-cybercrime-arup/ ↩

[5] Deloitte — *Generative AI is expected to magnify the risk of deepfakes and other fraud in banking*. https://www.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html ↩

[6] Sumsub — *Global deepfake incidents surge tenfold from 2022 to 2023*. https://sumsub.com/newsroom/sumsub-research-global-deepfake-incidents-surge-tenfold-from-2022-to-2023/ ↩

How Do We Get There? A Roadmap

"Wanderer, there is no road, / the road is made by walking."

— Antonio Machado, "Proverbios y cantares," XXIX (1912)

In 2035 the city of Barcelona stands as a living witness to a decision taken ten years earlier: in 2026, governments and companies decided to act. From the first moment the regulatory framework for proof of existence based on AI and blockchain was approved, trust infrastructures began to proliferate like filaments of a collective neural network. Today, along La Rambla, tourists do not just consult digital maps; they scan QR codes that, with a simple gesture, verify the origin and integrity of every product — from DO Penedès wine to art pieces made by 3D printing. Every label carries a small blockchain record that registers, in real time, the steps of production, transport and consumption, all while consuming barely any energy: the modern blockchain, after the Merge, has cut its energy cost by about 99.95% compared with the previous model, so anchoring the truth no longer squanders the planet [2].

On the facades of buildings, verifiable climate-impact data is projected: graphs showing how urban systems have optimized consumption thanks to artificial intelligence. These smart platforms adjust buildings' electricity use according to real demand and the availability of renewable energy, minimizing waste and maximizing efficiency. The investment needed for this change was estimated at between 1% and 2% of world GDP — a figure many critics considered too high at the time, but far cheaper than the cost of doing nothing — and the recovery of these costs has been swift, driven by operational savings and the market's new confidence [1].

In schools, students learn to use AI tools that help verify documents and news, making the spread of deepfakes harder. In a research center, a young researcher shows how a system trained to recognize digital manipulations can check a viral image against its origin in milliseconds, immediately sending the result to a public blockchain. This process — proof of existence, the heart of *t2p* — has become a pillar of digital trust: every piece of data circulating on the network is traceable back to the moment it was sealed, and any alteration is laid bare.

In the labor market, companies no longer depend solely on annual audits to validate certifications; many credentials are issued and verified instantly through smart contracts. A software engineer can prove, with a single transaction, that he has completed an open-source project, and his record remains verifiably registered for future collaborators.

This scenario is not a utopia; it is the possible culmination of a series of concrete decisions taken over a decade. Every actor — government, company, citizen — can adopt a set of clear principles, ranging from a commitment to standardizing proof of existence to choosing low-energy consensus protocols. The result is a society where trust is built on verifiable data, and where emerging technologies work as infrastructure for a sustainable, just future.

The three actors

Governments — lead the trust framework. The first step for a responsible government is to set a carbon price that reflects the real cost of climate damage and incentivizes decarbonization at every level of the economy. This price must be backed by international verification and monitoring (MRV) standards based on proof of existence. The modern blockchain, after the Merge, consumes far less energy than traditional blockchains, as we have seen, and can thus become the backbone of an immutable, verifiable registry of emissions data [2]. Governments can fund verification networks at national and regional scale, support research into ultra-efficient consensus protocols and pass legislation that steers companies toward ESG reports certified in a checkable way. This creates a chain of trust where every impact claim can be audited almost in real time, reducing the risk of greenwashing and increasing credibility with investors and citizens.

Companies — turn data into sustainable assets. For organizations, the key is investing in clean infrastructure that guarantees the traceability of their production processes. This means deploying blockchain solutions to anchor ESG data and using proofs of existence to validate every step of the product's life cycle, from raw-material extraction to final recycling. Companies that adopt this model not only comply with the emerging rules; they also gain a competitive edge: consumers and investors prefer brands that offer verifiable transparency. Adopting these technologies also cuts the costs of external audits

and traditional certifications, freeing resources for innovation and sustainable growth. The principle is simple: invest clean, anchor the data and do not greenwash.

Citizens — be vigilant and active. Every individual holds small but powerful levers: optimizing energy use at home, adopting electric mobility and verifying claims before sharing them. Blockchain-based tools let any user check the traceability of a piece of information in just a few clicks, preventing the spread of disinformation and reinforcing collective trust. Citizen participation in distributed monitoring programs — such as air-quality sensor networks or smart energy meters — also feeds publicly accessible data and serves as the foundation for more effective local policies. Combining personal action with decentralized verification tools creates an ecosystem where responsibility is distributed, increasing the pressure on governments and companies to act with integrity.

As we have seen, **governments** must set carbon prices and robust regulatory frameworks; **companies** must turn their ESG data into immutable, transparent assets; and **citizens** must exercise active vigilance, using proof-of-existence tools to validate information. This synergy multiplies the impact of every lever, turning the global challenge into a collective opportunity. It is time to get moving: **time2act**.

The trust infrastructure

The transition to a clean, resilient economy rests on a single premise: without reliable data there is no effective action. Throughout *time2act* we have shown how emerging technologies can transform uncertainty into assurance, and now is the moment to consolidate this thread. Artificial intelligence (AI) and the blockchain form a synergy that makes it possible to generate, validate and share information with a degree of integrity that was until now hard to reach. In this ecosystem, proof of existence — the central concept of *t2p* — acts as an immutable anchor for every kind of evidence: from a factory's carbon emissions to the green-energy certificates of a community microgrid.

1. Immutable data, agile decisions

The blockchain guarantees that every emissions record or ESG certificate remains unalterable and verifiable for any interested party. This reduces the need for costly, drawn-out audit processes, shortening validation time from weeks to minutes. This efficiency not only boosts competitiveness; it creates a solid base for public policies grounded in real evidence.

2. AI as the engine of analysis and prediction

AI algorithms process large volumes of data in real time, identifying hidden patterns and projecting future scenarios with ever higher precision. When these predictions rest on immutable records, their credibility grows: the results stop being mere estimates that are hard to check and become anchored in a verifiable layer of trust. This lets governments and companies take proactive measures — such as adjusting renewable energy production or redirecting investments — without the paralysis of doubt about the data's veracity.

3. Proof of existence: the seal of truth

t2p offers a technical solution for any information that requires authentication. Through a cryptographic hash and registration on a public blockchain, a proof of existence is generated that any actor — from the consumer to the regulator — can consult without intermediaries. This capability is crucial in the era of deepfakes and disinformation: every claim, certification or complaint can be checked against an immutable proof, reinforcing social and commercial trust.

4. Decentralized governance and citizen participation

The blockchain does not just store data; it also enables transparent governance mechanisms. Through smart contracts, organizations can set clear rules for managing digital assets, reward sustainable behavior and penalize infractions automatically. Citizens, moreover, have verification tools that let them exercise active vigilance over public policies and corporate practices, turning oversight into a collective, continuous process.

5. Route to action: principles for every actor

1. **Adopt low-energy blockchain platforms** — the Merge has cut modern chains' energy consumption by about 99.95% [2].
2. **Integrate AI for processing and analyzing ESG data**, ensuring every model feeds on verified records.
3. **Implement proof of existence as the validation standard** in certificates, reports and public statements.
4. **Foster citizen participation through open portals**, where any user can consult and validate the data with one click.
5. **Establish regulatory frameworks that recognize the blockchain as a source of verification**, simplifying audit processes.

As we have seen, trust is the bridge between the vision of a sustainable future and today's operational reality. The combination of AI and blockchain, reinforced by *t2p*, creates an infrastructure that turns data into truth and truth into action. The challenge is no longer technological but collective; it falls to every company, government and citizen to get moving. It is time to act with resolve and clear convictions: **time2act**.

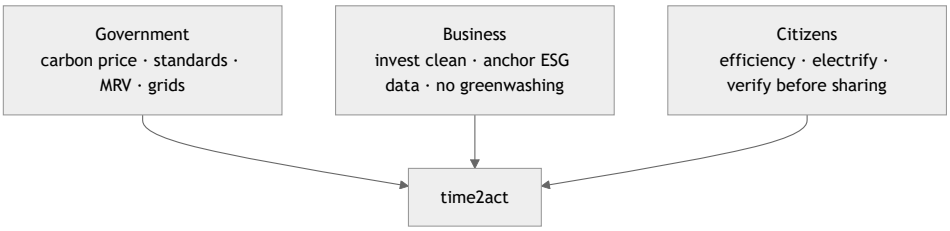


Figure 111.1 — Each actor — government, company and citizens — has concrete actions for the short, medium and long term.

The principles of action

1. Electrify everything possible. The energy transition is no longer an option, but the base on which any sustainable model is built. Every piece of industrial equipment, every transport fleet and every urban infrastructure should be powered by electricity, preferably generated from renewable sources. The key is to set a scalable migration calendar: identify the critical points, replace equipment with high-efficiency motors and seize smart-charging opportunities to balance demand.

2. Do it with clean, firm energy. Electrifying is not enough if the electricity comes from fossil origins. Combining AI to optimize energy spending and blockchain to record in real time the provenance of the current helps guarantee a truly clean supply chain. When every kilowatt-hour can be certified through a smart contract, clean energy stops being just an environmental aspiration and becomes a lever of competitiveness and consumer trust as well.

3. Measure and verify incorruptibly. The truth of the data becomes the new most valuable asset. Through *t2p*, every record of energy consumed, of emissions or of operational results can be encoded on an immutable blockchain, making manipulation difficult and enabling much faster audits. Precise measurement and instant verifiability are, therefore, the pillars that sustain any strategy of real impact.

4. Align incentives. The price of carbon must reflect its social cost. Putting a price on emissions and removing fossil-fuel subsidies creates an economic framework where green decisions are also the most profitable ones. This alignment incentivizes responsible behavior and steers investment toward clean technologies, without the need for excessive regulatory intervention.

5. Protect truth as a public good. In an era of deepfakes and disinformation, trust becomes a scarce resource. Blockchain-based infrastructure, reinforced by *t2p*, makes critical information on energy, emissions and social impact immutable and accessible to all interested parties. Public policies must recognize truth as a common good, establishing legal frameworks that protect certified data and deter manipulation. That way, every citizen, company or government can rely on a verifiable source to make informed decisions.

These five principles are not just theoretical recommendations; they are concrete steps that any actor — from the nimblest start-up to the national government — can begin to implement today. By following them, we turn the potential of AI and blockchain into a trust infrastructure capable of transforming data into truth and truth into action. **time2act**.

The regulatory horizon

No roadmap would be complete without looking at the legal framework now approaching, because regulation will reward those who move early. In Europe, Article 50 of the Artificial Intelligence Regulation (the EU AI Act) will require machine-readable marking of AI-generated content and the labeling of *deepfakes*; its obligations apply from August 2, 2026 [3]. In parallel, the eIDAS 2.0 regulation rolls out the European digital identity wallet and introduces new figures — qualified electronic ledgers and qualified electronic archiving — that give legal recognition to the immutable audit-trail model this book defends [4].

Across the Atlantic, the United States has passed the TAKE IT DOWN Act (signed in May 2025), which criminalizes non-consensual intimate imagery — including AI-generated imagery — and requires platforms to remove it within 48 hours; and the DEFIANCE Act, passed by the Senate in January 2026, adds a civil avenue for victims to claim redress [5].

The message for governments and companies is unequivocal: the legal expectation is no longer just to "keep" information, but to be able to PROVE its origin and integrity. Organizations that adopt proof of existence today — anchoring the hash of every critical piece of data before sharing it — will not comply with the rules grudgingly and at the last minute; they will arrive prepared, because they will already have made verifiability a habit. Regulation, then, is not an added burden on the transition: it is the same road, signposted.

The promise and the peril (synthesis)

For the first time in contemporary history we have a clear diagnosis, proven tools and an infrastructure capable of protecting the truth that underpins our decisions. Artificial intelligence, fed by large volumes of clean, verified data, can transform scattered information into actionable knowledge with a precision that decades ago was impossible. In parallel, the blockchain — in its most efficient form after the Merge, which has cut energy consumption by about 99.95% — offers a layer of impartiality and immutability that turns every transaction, every intellectual-property record or every climate certification into a proof of existence that is hard to forge [2]. This combination creates a trust infrastructure where truth is not only detected but kept intact against attempts at manipulation, such as deepfakes and mass disinformation campaigns.

Yet the story of the opportunity is inevitably entwined with an imminent danger. Acting has a cost — between 1% and 2% of world GDP to hold warming below 2 °C — but it is far cheaper than the cost of not acting, which grows with every year of delay [1]. Moreover, technology, however powerful, is no panacea. The blockchain can guarantee the integrity of records, but without political and business will determined to adopt it, those records lie dormant, like files in a locked warehouse. Similarly, AI models can offer useful predictions only if reliable data is available, along with governance processes that ensure their ethical, transparent use.

The window for seizing this synergy is limited. Every year that passes without coordinated action raises the economic cost and complicates the technological path, because legal and regulatory infrastructures fall behind the pace of innovation. The barriers to adoption — from the lack of specialized talent to the opposition of entrenched interests — sink us once the key moment disappears. That is why a collective will is needed to drive public policies, strategic investments and cross-sector collaborations to deploy this trust infrastructure at global scale.

We have the tools and the conceptual framework needed for truth to be the foundation of every decision; but if we do not act now, the door will close, the costs will rise and our capacity to act on reliable information will be compromised. The key is to combine technological power with a will determined to turn this promise into concrete reality. It is the moment to say **time2act**.

This vision is reinforced by Beth Noveck, who after the Horizon scandal — where more than 900 unjust convictions led to at least thirteen suicides — proposes that artificial intelligence be a civic infrastructure, **verifiable, auditable and accountable to the people it impacts** [6]. Her procurement framework obliges governments and companies to demand records anchored on a neutral blockchain, independent of the creator, ensuring that every AI model can be reviewed and reproduced without black boxes. This requirement makes proof of existence a minimum condition for any public tender, transforming data and algorithms into open assets that can be validated in real time, and it completes our roadmap: we do not just protect the truth — we make its transparency mandatory from the first moment of procurement.

It's time2act

Twenty years ago, Paul Roberts, in *The End of Oil* (2004), painted the picture of a world running out of easy energy, with a fragmented geopolitics and a climate showing its first irreversible cracks. His diagnosis, as precise as it was alarming, still echoes in every boardroom and every university lecture hall. But the landscape has changed: today we have not only the knowledge of the risks but also a technological infrastructure capable of restoring the trust the crisis has eroded.

Principle 1 — Truth as the operating base. Proof of existence turns any piece of data — from energy-origin certificates to climate audit results — into an immutable, verifiable record. This reduces the uncertainty that feeds deepfakes and market manipulation, and enables decisions based on checkable facts.

Principle 2 — Energy efficiency as a requirement. As we have seen, the modern blockchain consumes far less energy than the first iterations [2]. This means the trust infrastructure does not become a new source of runaway demand, but a tool that can coexist with the systems already in place.

Principle 3 — A sustainable cost of adoption. Holding warming below 2 °C requires investing between 1% and 2% of world GDP, a figure that pays for itself many times over when compared with the cost of inaction: more climate disasters, lost trust in the markets and rising raw-material prices [1]. The transition is not free, but it is far cheaper than its opposite.

Principle 4 — Scalability for every actor. From national governments to renewable-energy start-ups, the combination of AI for data processing and blockchain for its certification creates a layer that can be adopted progressively. Each entity can start a local pilot — for example, validating the origin of energy in an industrial zone — and scale toward broader supply chains.

Principle 5 — Participatory governance. The transparency inherent in these systems makes it possible to include citizens, academics and NGOs in verification and audit processes, reinforcing the social contract and generating an ecosystem where trust is not a concession but a natural consequence.

The path opening before us is clear: integrate AI and blockchain as pillars of a new architecture of trust, deploying proof of existence to guarantee that every decision rests on authentic data. It is not enough to roll out technology; we must transform mindsets, establish agile regulatory frameworks and create cross-sector collaborations that catalyze change.

So the call is not for more diagnoses or empty promises, but for setting in motion the tools we already have. The time for talking has passed; now it is time to act with resolve and precision. **It's time2act**.

---

[1] (cost of acting: 1-2% of GDP) Oxford Economics / CPI. https://www.climatepolicyinitiative.org/the-cost-of-inaction/ ↩

[2] ethereum.org, *The Merge*. https://ethereum.org/roadmap/merge/ ↩

[3] EU AI Act, Article 50 (transparency of AI-generated content; applies 2-Aug-2026). https://artificialintelligenceact.eu/article/50/ ↩

[4] Signaturit — *eIDAS 2.0: qualified electronic ledgers and archiving*. https://www.signaturit.com/blog/eidas-2-regulation/ ↩

[5] Orrick — *TAKE IT DOWN Act becomes law*; DEFIANCE Act (Senate, Jan-2026). https://www.orrick.com/en/Insights/2025/05/TAKE-IT-DOWN-Act-Becomes-Law ↩

[6] Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 8; Post Office Horizon Inquiry (Sir Wyn Williams, 2024). ↩

The Seal and the Chain: Who Controls the Proof, and Who Gets Trapped

"Power tends to corrupt, and absolute power corrupts absolutely."

— Lord Acton, letter to Mandell Creighton (1887)

At the entrance to an improvised laboratory in Nairobi, a line of people waited to be inspected. Among them, young students, street vendors and families with small children lined up under a poster announcing "World ID — Prove you are human." A team of engineers, wearing protective masks and a high-resolution camera, scanned each person's iris while a monitor displayed numbers that seemed to confirm their identity. In less than two minutes, the system recorded an immutable "proof" on a public blockchain, assigning them a cryptographic credit as a reward for taking part.

This scene, which could appear in any optimistic report on the fight against bots and deepfakes, is exactly what *time2act* has defended: a reliable mechanism to prove that behind an avatar there is a real person. Yet the irony becomes palpable when, only two years later, the same Worldcoin/World ID project faces bans and fines in more than nine jurisdictions. Kenya declared it unconstitutional, and both Kenya and Colombia ordered the total destruction of the biometric data collected. What had begun as a solution to reinforce digital trust turns into a clear example of the danger lurking in any immutable infrastructure [1].

This paradox is the starting point of this chapter. The infrastructure that makes t2p possible — biometrics, immutable registry and blockchain — can be, simultaneously, the architecture of mass surveillance. This is not a rejection of the concept; the mechanism is necessary to fight digital forgery. What we must recognize, as Snowden establishes, is that "the credo of the mass surveillance system was 'collect it all, process it all, exploit it all, partner it all,'" and that this pattern has already left an indelible mark on our recent history [20].

This chapter presents itself as the book's honest counterweight. Following the mandate of *time2act* — "the promise and the peril" — it gathers, in one place, the criticisms that until now had been scattered across ethics reports, legal analyses and whistleblower testimony. The central thesis is clear: technical decentralization does not guarantee democratic governance; the real problem is the power relations that form around control of the proof. Who defines the rules of entry to this registry? Who can delete or modify the data without accountability? Which interests benefit when a single entity controls both access and verification?

In the current context, where the big Ethereum L2s operate with a single corporate sequencer and 42% of blocks are censored according to the OFAC lists [5], the question left open is who verifies the verifier. The blockchain can say "don't trust, verify," but without independent, transparent auditing, this verification becomes a subtle form of control, as Dominic Frisby reminds us: "the blockchain's motto is 'don't trust, verify'; but the question left open is who verifies the verifier" [32].

Thus this opening chapter highlights that, despite the promises of restored trust, the same technology can be used to consolidate power and impose a surveillance that limits individual freedom. The answer is not to dismiss biometrics or the immutable chain, but to build frameworks of democratic governance that ensure the "seal" serves to protect us and not to trap us. The reader's next step will be to confront these tensions and consider how to guarantee that control of the proof remains at the service of society, and not in the hands of a technological elite.

The neutrality that does not exist

At the close of 2026, the infrastructure sustaining the Ethereum network is dominated by a single entity: a corporate sequencer run by Coinbase through Base. This is not just a matter of performance; this single node decides which block gets included and into which confirmation chain it passes, exercising a form of censorship that affects 42% of Ethereum blocks via relays that apply the OFAC lists. Moreover, the second validation layer that t2p uses today has concentrated block production in a small group of validators chosen after the "Rio" upgrade (2025) [5].

This concentration contradicts the popular myth that the blockchain is inherently decentralized and neutral. The truth is that whoever programs and operates the nodes determines which records are kept and who has access to them. As Cathy O'Neil points out, the mathematical models sustaining these

infrastructures are "opaque and invisible to everyone except the high priests of the sector: the mathematicians and the computer engineers" [4]. Her observation reveals that trust does not rest on the code, but on the people who write and maintain it.

Marc Goodman sums it up clearly: "whoever controls the code will control the world, for better and for worse" [3]. This is no hyperbole; the ability to define the rules of an immutable chain translates into power over digital identity, intellectual-property rights and even the capacity to exclude political or social actors. Edward Snowden confirms this dynamic: "these machines do what they do because someone orders them to, in a very concrete and painstaking way; and that someone can even be a seven-year-old child" [2].

At the table where the rules of the game are drafted, the margin for breaking them is enormous. Jack Watling observes that "being at the table where the rules of the game are set increases a state's margin for breaking those same rules to its own advantage" [6]. This power is not limited to states; private infrastructures like Starlink, controlled by corporations without democratic accountability, have become strategic assets that shape global communication and, therefore, the capacity for surveillance and control [7].

The historical parallel is revealing. Dominic Frisby recalls how Isaac Newton, as sole head of the royal mint (Master of the Mint), concentrated control over the validation of the currency and pursued counterfeiters with implacable rigor [8]. The centralization of the power of verification is not a new phenomenon; it has merely changed from metal to bits.

Without an explicit democratic governance that regulates who can be a validator, what their obligations are and how decisions are supervised, proof of existence (t2p) runs the risk of transferring power from the traditional intermediaries — banks, notaries, public authorities — to a technical elite operating in the shadows; as we have seen. The immutable infrastructure thus becomes a new kind of intermediary: invisible but decisively influential, and without accountability mechanisms to guarantee its neutrality.

Crystallized consent

Within the framework of the "verified consent" that underpins proof of existence, the book has presented this seal as a guarantee of transparency and responsibility. Yet the mere existence of a recorded acceptance does not ensure that it was autonomous or free. Carissa Véliz reminds us that "you have no margin to negotiate the terms and conditions, which, moreover, can change at any moment without prior notice" [9]; and the opinion of the European Data Protection Board stresses that the "consent or pay" model does not constitute "freely given" consent [10]. This means that when the blockchain permanently seals that "I accept," it transforms a possibly voluntary choice into an immutable and therefore inescapable fact.

Two concrete examples illustrate how this dynamic materializes in registry infrastructures. At the Casa de Contratación in Seville, the registry of obligations was imposed through mandatory oaths and penalties that reached two hundred lashes for breach [11]. That "immutable registry" served as a tool of coercion: the act of signing was not a free declaration but a legal requirement which, if refused, carried severe economic and personal penalties. In this context, the blockchain could consolidate the memory of those oaths, turning them into eternal witnesses of institutional pressure.

Another vector of coercion comes from the energy dependence that sustains decentralized networks. Jack Watling describes this situation as a "finger pressed on the throat" [12], a condition that invalidates the autonomy of consent because operators cannot dissociate themselves from critical infrastructure without compromising its operation. The need for a constant energy supply for nodes and validators creates a structural vulnerability: the "click" that attests consent becomes a forced act, driven by the fear of losing access to essential resources.

These cases reveal a fundamental contradiction in the t2p proposal. If the technology records and verifies any acceptance without discerning between will and imposition, the result is a form of digital authoritarianism that reinforces existing power hierarchies. The solution is not just to add more layers of cryptographic security; governance mechanisms must be introduced that can validate the legitimacy of the consent itself. This means, for example, establishing independent audit protocols that review contractual conditions and ensure the registry does not become a tool of structural coercion.

Ultimately, the chain can be a reliable witness only if its testimony is genuine. Clear distinctions between real and forced consent must be built into the design of t2p; otherwise, the eternal "click" that gets

recorded could end up perpetuating the very forms of control it aims to fight. The challenge, then, is not only technological but institutional: the verification of consent must be accompanied by a democratic framework that protects individuals' autonomy and prevents the seal from becoming a chain of surveillance.

The surplus value of the proof

When we tokenize human experience as "proof of existence," what is apparently a tool for restoring trust becomes, in reality, a tradable asset. Shoshana Zuboff describes this phenomenon as surveillance capitalism: "human experience [is] free raw material" that platforms extract, predict and sell [13]. Every time a biometric fingerprint or a timestamp is recorded on an immutable chain, a new unit of value is added to the database that feeds the market in personal information. Zuboff herself reminds us that "we are the ones being hoarded; our experience is the object of extraction" [14], underlining that control of the data flow translates into economic and political power.

This dynamic is not merely theoretical. In 2024, the United States Federal Trade Commission banned data brokers for the first time from selling sensitive locations and fined Avast $16.5 million [15]. The same year, the Dutch Data Protection Authority (DPA) imposed a €30.5 million penalty on Clearview AI for having built a biometric database of thirty billion photographs scraped without consent. These actions show that the commodification of the immutable registry has already generated a revenue model based on exploiting proofs of identity, and that regulatory mechanisms are still struggling to keep pace.

Dominic Frisby's historical warning about the silent debasement of standards offers a long-term perspective. The Roman denarius, according to Frisby, went from 95-98% silver to less than 1% without anyone noticing [16]. This decline happened because there was no independent, continuous audit watching over the quality and integrity of the monetary system. In parallel fashion, immutable-registry technologies can suffer a similar erosion if their operation is left subject to private interests without public oversight. The lack of a permanent audit leaves the door open to subtle manipulations: changes in consensus algorithms, the insertion of nodes controlled by corporate entities and, finally, a concentration of power that contradicts the decentralized ideal.

Concrete safeguards must therefore be established against the commodification of recorded data. First, the protocol must include external verification mechanisms, with independent auditors accredited by democratic institutions and not by companies that profit from them. Second, a legal framework must be put in place that mandates transparency about the owner of the seals and the associated chains, guaranteeing that any transfer of value is public and reversible. Finally, a system of participatory governance is needed in which users can exercise control over the use policies and the fees applied to their own proof.

Only with these measures can we prevent the "seal" from becoming a "chain" of exploitation. The surplus value of the proof lies not only in its technical capacity to validate authenticity, but in how society regulates who profits from it and what limits are placed on its use. Without a permanent audit and democratic governance, the risk is that, just as happened with ancient coinage, the technical standards will silently degrade until they lose their protective function entirely.

The right to be forgotten, deleted

The moment a piece of information is engraved on an immutable chain, its existence is no longer subject to the whims of institutions or the passage of time. Immutability, celebrated as a guarantee of trust, simultaneously becomes the anchor of a permanent tracking that forecloses any possibility of forgetting. Edward Snowden describes this phenomenon with unsettling precision: "I found it atrocious to be so inextricably, technologically bound to a past I regretted completely but could barely remember" [20]. This is not just a personal reflection; it is the reality of a system designed to "collect it all, process it all, exploit it all, partner it all," and at the center of whose operations the Bluffdale data center was built to store "the pattern of life of the entire planet" [17].

These processes go beyond a simple accumulation of data: the network has imposed a "fidelity to memory, identitarian uniformity and, therefore, ideological conformity" [18]. When the registry is immutable, any error, obsolete datum or piece of information that may someday prove harmful to the subject becomes a permanent wound. There is no mechanism of deletion or annulment; digital memory remains alive and accessible indefinitely.

As already noted in 2026, this tracking power has materialized in the massive use of chain-analysis tools by US federal agencies. ICE, the IRS and the FBI have acquired solutions such as Chainalysis or TRM Labs with the explicit aim of de-anonymizing the immutable registry, linking digital addresses to specific people [19]. These technologies make it possible to follow the flow of digital assets, identify transactions and, finally, trace individuals back to their biological or geographic origin. The result is a form of surveillance that transcends the traditional boundaries of state control: the information no longer depends solely on a user's consent, but becomes an immutable witness to every action.

Frisby stresses that even gold, considered the "only money that is not a promise," becomes vulnerable when its ownership and location are fully transparent: "an easy target for confiscation or taxation" [21]. Analogously, an immutable registry becomes an asset that any authority can exploit to impose sanctions, confiscate property or restrict civil rights. The impossibility of deleting or revoking erroneous data means citizens have no effective recourse against unjust decisions based on inaccurate information.

This situation poses a fundamental contradiction: immutability, which the book has celebrated as a pillar of trust, also becomes the instrument that removes any possibility of forgetting and, therefore, prevents the rectification or deletion of harmful data. Without a permanent audit and democratic governance, this immutable registry turns against the citizen, becoming a tool of coercion rather than a guarantee of transparency. The challenge, then, is not only technological but political: mechanisms of revocability and expiry must be established for erroneous or transitory records, ensuring that immutability does not translate into a stigma of permanent injustice.

Unappealable algorithms

The moment an entity decides to seal the provenance of a piece of data, the criterion behind that seal is hidden behind a wall of code and mathematics incomprehensible to most. As Cathy O'Neil warns, the verdicts issued by these models are "unappealable" and, at the same time, "opinions embedded in mathematics" that only the high priests of the sector — engineers and data scientists — can read and question [22]. This opacity turns the seal into an unchallengeable authority and creates a power that exceeds any democratic control.

Carissa Véliz reinforces this perspective: predicting human behavior amounts to trying to control it, a dynamic that leads inevitably toward authoritarianism when the predictions are used to regulate or restrict freedoms [23]. The model that guarantees the "veracity" of an immutable record can be used to impose limits on individuals without their knowledge or capacity to contest them. Edward Snowden adds: "predictability is manipulation." When a system anticipates future actions, it does not merely offer information; it creates a subtle mechanism of coercion that conditions present decisions to fit the pre-established pattern [24]. The seal thus stops being a simple guarantee of provenance and becomes an instrument of conformity.

This concentration of algorithmic power is reinforced by the fact that whoever sets the rules of the game has the margin to break them. Jack Watling notes that "whoever makes the rules has margin to break those same rules to their own advantage" [6]. In an ecosystem where the blockchain infrastructure and the immutable registry are controlled by a few corporate actors, this capacity translates into a structural vulnerability: the seal can be modified or ignored according to the interests of those who created it.

As we have seen, the legal framework of 2026 retreats before these threats; the postponement of the obligations for high-risk AI within the "Digital Omnibus" package and the repeal of Colorado's AI law leave a regulatory vacuum that facilitates the indiscriminate expansion of sealing systems [25]. Without clear rules imposing transparency and independent review, unappealable algorithms can operate without brake or counterweight.

The seal's fragility is also evident in the field of geolocation. During the war in Ukraine, electronic interference spoofed GPS signals, proving that neither time nor position is invulnerable to deliberate manipulation [26]. This reveals a fundamental truth: the binding force of the seal is only as strong as the reliability of the underlying infrastructures. If these can be falsified, mechanisms of verifiable explainability and rectification must be developed with the same legal robustness as the original immutable record.

In short, the concentration of algorithmic power turns the seal into a potential weapon of control. The solution lies not only in improving the technology, but in establishing a democratic governance capable of

imposing transparency, independent review and rectification mechanisms that prevent the "code" from becoming the sole guardian of the future.

The self-censorship of the permanent record

In the spring of 2024, a young Catalan writer withdrew an essay on artificial-intelligence practices during the election period, fearful that the mere act of publishing it might attract state surveillance. This behavior is not anecdotal; according to Carissa Véliz, between a third and two thirds of writers avoid certain topics for fear of government surveillance [27]. The deep cause lies in what scholars call *mission creep*: an infrastructure designed for one purpose — like the t2p proof of existence — progressively extends toward unforeseen uses, transforming itself into a weapon of social control.

The immutable registry, the central element of any blockchain that aims to guarantee integrity and traceability, thus becomes a "seal" that can be exploited to impose functional and temporal limits on freedom of expression. A recent example is Eurodac, the European fingerprint system for asylum procedures. In 2024 the minimum age for biometric collection was lowered from twelve to six years, opening the database to expanded police access and enabling early surveillance of minors [28]. This expansion is not merely technical; it is a political decision that transforms a migration-management tool into a source of information for social control.

The proliferation of the permanent record also feeds more subtle but equally devastating chilling effects. In Syria, during the war, disinformation orchestrated by Russian actors created a fabricated uncertainty that paralyzed international and local responses [29]. When knowledge becomes uncertain, the capacity to act autonomously evaporates: people begin to self-censor to avoid becoming the target of possible sanctions or persecution. This phenomenon is reflected in Shoshana Zuboff's observation that the system's fundamental power is to determine "who knows, who decides, and who decides who decides" [30]. In this scenario, the immutable registry stops being a guarantee of transparency and becomes a mechanism that consolidates information in the service of a few authorities.

The legitimacy of this power comes into question when platforms become "the real arbiters of what people can say," as Véliz notes: we must ask whether they possess the right kind of legitimacy to exercise this control [31]. "Protecting" biometric data or digital transactions is often presented as a euphemism for "controlling" individuals' behavior and identities. The direct consequence is a self-censorship that extends beyond speech, limiting creative initiative, inhibiting critical research and encouraging self-exclusion from uncomfortable topics.

To prevent a tool built for one purpose from becoming the infrastructure of another, explicit limits — both functional and temporal — must be established. This means clearly defining which data can be recorded, for how long and under what independent supervision. Only with transparent review mechanisms and democratic participation can we guarantee that the digital seal does not become a chain of oppression. In this delicate balance between trust and surveillance, the real danger lies not so much in the technology itself as in who controls it and how far its reach is allowed to expand.

The counterargument t2p must survive

Across the long journey of this book we have shown how proof of existence can restore trust in a world saturated with deepfakes and digital forgeries. Now we face the most uncomfortable reflection: the same infrastructure — biometrics, immutable registry and blockchain — is the skeleton of a surveillance architecture that could turn the digital seal into a chain of oppression. The key to not losing ourselves in a techno-utopian scenario is recognizing that **technical decentralization does not equal democratic governance**; the fundamental problem is the power relations that take shape around control and oversight.

The chain's motto, "don't trust, verify," has become a universal axiom of distributed security. Yet, as Dominic Frisby points out, the open question is **who verifies the verifier?** [32] If the verification process stays inside a black box operated by a few actors, what should be a guarantee of transparency turns into an instrument of coercion. **As we have seen, in 2026 42% of blocks are built through relays that apply OFAC censorship**; the chain t2p uses today concentrates block production in a small group of chosen validators. [5] This concentration reveals how, despite an architecture decentralized at the protocol level, effective control rests in corporate and state hands.

For **t2p** to fulfill its mission without becoming an instrument of surveillance, four essential demands must be established:

1. **Democratic governance of the nodes** — Whoever operates, validates and audits the nodes must be subject to open, participatory processes, avoiding the concentration observed in the L2 chains. Oversight must come from independent, representative entities, not only from large financial platforms.

2. **Transparent auditing of the algorithms** — Verification methods must be subject to periodic review by external experts, with results published in accessible formats. This counters the opacity described by Cathy O'Neil, whereby models are "opinions embedded in mathematics," readable only by the "high priests of the sector." [4]

3. **A mechanism of revocability and correction** — When an immutable record contains errors or data collected without consent (like Eurodac, which has lowered the biometric age to six years and allows police access), a safe path for rectification or deletion must be included, guaranteeing the right to be forgotten and preventing the coercive use of erroneous data. [28]

4. **Control over the surplus value of the data** — It must be clearly defined who captures and commercializes the information derived from the registry; without this control, a surveillance economy is perpetuated in which personal information becomes a high-yield commodity for corporations and states. [15]

In short, **the chain seals probabilities, not truths**; an immutable seal can anchor coercion as effectively as authenticity. Recognizing this duality does not weaken the message of *time2act*; it armors it: only by confronting the limits of power and establishing strict rules of governance can we guarantee that proof of existence is a tool of trust and not an instrument of oppression.

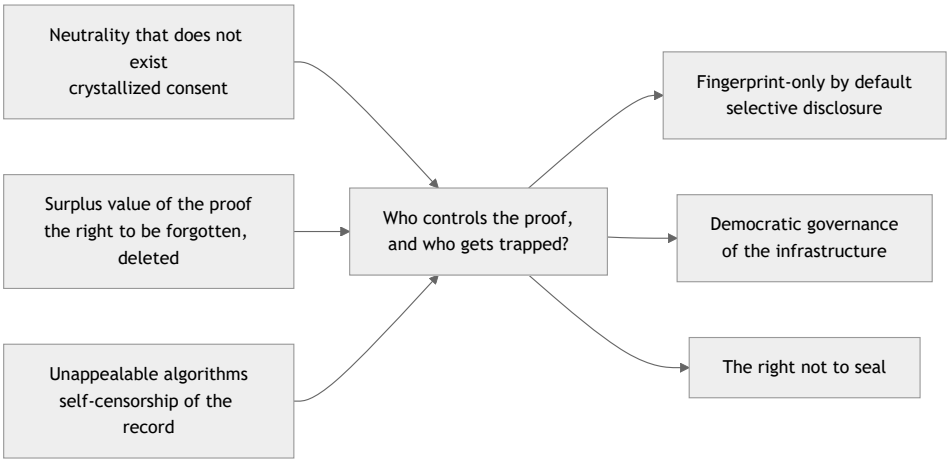


Figure 112.1 — The seal and the chain: who controls the proof, and who gets trapped. The risks the chapter unpacks — the neutrality that does not exist, crystallized consent, the surplus value of the proof, the right to be forgotten deleted, unappealable algorithms and the self-censorship of the permanent record — and the counterweights it demands: hash-only by default, selective disclosure, democratic governance of the infrastructure and the right not to seal.

It's time to act

The proof of existence we have built throughout *time2act* is, like any seal, a double-edged tool: it authenticates the truth and, simultaneously, can become the anchor of a new form of coercion. For this instrument to remain in the service of trust and not of control, five design principles must guide its governance.

1. Transparent, distributed governance of the nodes. Technical decentralization does not by itself guarantee digital democracy; power remains concentrated in whoever controls the points of entry. As we have seen, in 2026 42% of Ethereum blocks are built through relays that impose OFAC censorship, and the L2 networks operate with a single corporate sequencer such as Coinbase's Base [5]. The nodes that record the proof must therefore be subject to continuous public audit and to selection based on open criteria (technical competence, geographic diversity and commitment to democratic values). Transparency

must be verified by independent entities, as Frisby reminds us: "don't trust, verify" — but the key question is *who* verifies the verifier [32].

2. The right to revocation and expiry. An immutable record cannot be eternal if it aims to respect people's fundamental rights. European legislation has moved toward the "consent or pay" model, but without a mechanism that allows consent to be withdrawn, this model becomes a form of captivity [10]. Protocols must be established that make it possible to revoke the identity certificate and to have biometric records expire automatically when they are no longer necessary or become vulnerable to abuse.

3. Real biometric minimization. The centralization of biometric databases, like the Eurodac project, which has lowered the collection age to six years and shares the data with police forces [28], shows how an infrastructure conceived for security can be turned into a tool of mass surveillance. The solution is to avoid centralized databases and use instead decentralized proofs based on "zero-knowledge" techniques that reveal no sensitive data yet allow an individual's existence to be validated.

4. Independent algorithmic auditing. The mathematical models behind automated decisions are, according to O'Neil, opaque and accessible only to "the high priests of the sector" [4]. To guarantee that the proof does not become an instrument of injustice, independent commissions must be instituted with legal power to review and validate the algorithms that process biometric and blockchain data.

5. End-to-end encryption as the real defense. Snowden affirms that end-to-end encryption is the only effective protection against surveillance, since it turns intercepted data into incomprehensible noise without the keys [34]. Every transmission of proofs must be protected by this level of encryption, and the keys must be managed in a decentralized way, avoiding single points of failure.

These principles set functional limits against the *mission creep* that, as Zuboff warns, transforms human experience into free raw material for extraction and sale [13]. Only if democratic controls are installed — independent audits, revocation of consent and a transparent architecture — can proof of existence remain an emancipatory tool rather than an instrument of capture. **The seal can guarantee the truth; the chain must prevent that truth from becoming a cage.**

---

[1] Biometric Update / ICLG — *Worldcoin/World ID: bans and deletion orders (Kenya unconstitutional; Colombia)* (2024-2025). https://www.biometricupdate.com/202510/colombia-orders-world-shut-down-citing-biometrics-compliance-failures ↩

[2] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

[3] Marc Goodman — *Delitos del futuro* (Ariel, 2015). ↩

[4] Cathy O'Neil — *Armas de destrucción matemática* (Capitán Swing, 2018). ↩

[5] L2BEAT — *Decentralization status of Ethereum L2s (sequencers)* (2026). https://l2beat.com/scaling/summary — and MEV Watch, *blocks with OFAC censorship* (2026). https://www.mevwatch.info/ ↩

[6] Jack Watling — *Statecraft: The New Rules of Power in a Divided World* (2023). ↩

[7] Jack Watling — *Statecraft: The New Rules of Power in a Divided World* (2023). ↩

[8] Dominic Frisby — *The Secret History of Gold* (2024). ↩

[9] Carissa Véliz — *Privacidad es poder* (Debate, 2021). ↩

[10] EDPB — *Opinion 08/2024 on "consent or pay"* (17-04-2024). https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-082024-valid-consent-context-consent-or_en ↩

[11] Dominic Frisby — *The Secret History of Gold* (2024). ↩

[12] Jack Watling — *Statecraft: The New Rules of Power in a Divided World* (2023). ↩

[13] Shoshana Zuboff — *La era del capitalismo de la vigilancia* (Paidós, 2020). ↩

[14] Shoshana Zuboff — *La era del capitalismo de la vigilancia* (Paidós, 2020). ↩

[15] FTC — *Actions against data brokers (X-Mode/Outlogic) and Avast* (2024). https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-order-prohibits-data-broker-x-mode-social-outlogic-selling-sensitive-location-data — and Autoriteit Persoonsgegevens (NL), *€30.5 million fine on Clearview AI* (2024). https://www.edpb.europa.eu/news/national-news/2024/dutch-supervisory-authority-imposes-fine-clearview-because-illegal-data_en ↩

[16] Dominic Frisby — *The Secret History of Gold* (2024). ↩

[17] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

[18] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

[19] FedScoop — *US federal agencies and blockchain analytics (Chainalysis, TRM Labs)* (2024-2025). https://fedscoop.com/ice-wants-more-blockchain-analytics-tech/ ↩

[20] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

[21] Dominic Frisby — *The Secret History of Gold* (2024). ↩

[22] Cathy O'Neil — *Armas de destrucción matemática* (Capitán Swing, 2018). ↩

[23] Carissa Véliz — *Prophecy: Prediction, Power, and the Fight for the Future* (2025). ↩

[24] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

[25] Council of the EU / Pinsent Masons — *"Digital Omnibus": postponement of the high-risk AI rules* (May 2026). https://www.pinsentmasons.com/out-law/news/rules-high-risk-ai-delayed-under-eu-omnibus-deal — and Colorado General Assembly, *SB 26-189* (2026). https://leg.colorado.gov/bills/sb26-189 ↩

[26] Jack Watling — *Statecraft: The New Rules of Power in a Divided World* (2023). ↩

[27] Carissa Véliz — *The Ethics of Privacy and Surveillance* (Oxford University Press, 2024). ↩

[28] European Policy Centre — *Eurodac expansion (age 12→6; police access), 2024 Migration Pact*. https://www.epc.eu/publication/managing-migration-through-data-the-expansion-and-integration-of-eurodac/ ↩

[29] Jack Watling — *Statecraft: The New Rules of Power in a Divided World* (2023). ↩

[30] Shoshana Zuboff — *La era del capitalismo de la vigilancia* (Paidós, 2020). ↩

[31] Carissa Véliz (ed.) — *The Oxford Handbook of Digital Ethics* (Oxford University Press, 2022). ↩

[32] Dominic Frisby — *The Secret History of Gold* (2024). ↩

[33] time2act — *Style guide (STYLE.md): "the promise and the peril"; "no techno-utopia without evidence"*.

[34] Edward Snowden — *Vigilancia permanente* (Planeta, 2019). ↩

The Siege on Truth: The Witness Who Cannot Be Silenced

"Our mission is to report these horrors of war with accuracy and without prejudice."

— Marie Colvin, "Truth at all costs," St Bride's, London (2010)

Toward ten in the morning on August 25, 2025, the Reuters camera mounted at the top of an external staircase of Nasser Hospital, in Khan Younis, in southern Gaza — the spot television crews had been using for days for its view of the city — was broadcasting live, as it did every morning. The first missile struck exactly there. The signal cut out dead: Hussam al-Masri, the Reuters cameraman who was filming, died with his lens. Some ten minutes later, when rescuers and journalists had already climbed the stairs to look for the wounded, the second strike arrived — the tactic known as the "double tap" [3]. The toll: at least 22 dead, among them five journalists — al-Masri (Reuters), Mariam Dagga (an Associated Press contributor), Mohammed Salama (Al Jazeera), Moaz Abu Taha (a Reuters contributor) and Ahmed Abu Aziz (Middle East Eye and Quds News Network) [1][3]. Later analyses of the videos, by NBC and by CNN, counted up to four munitions [3]. Prime Minister Netanyahu's office called it "a tragic mishap," and the army announced an investigation [3].

There is an atrocious irony in this scene, and it deserves to be looked at head-on: the proof of the attack was filmed by the very camera the attack was destroying. The Reuters live feed is at once the last image of Hussam al-Masri and the evidence of how he died. In this war, the image no longer accompanies the facts: it is the battlefield.

The scene takes place, moreover, in a war the world has not been allowed to go and see: since October 2023, the international press has been barred from entering Gaza independently [6]. And it takes place in the conflict the Committee to Protect Journalists (CPJ) describes as "the deadliest and most deliberate effort to kill and silence journalists" it has ever documented [1]. Closing the door to those outside, killing or marking those inside, and denying the images that get out despite everything — "staged," "actors," "this is AI" —: three strategies converging on a single goal, that no proof survive. This entire book has argued that trust can be made verifiable; this penultimate chapter takes that argument where everything is hardest and most necessary at once: where reporting kills and where denying is free. The thesis is the same as ever, taken to the limit: if every photo and every video is born with a proof of existence anchored on a public chain that no state, no platform and no consortium controls, killing the witness stops killing the proof — and denying it stops being free.

This chapter walks on stones the book has already laid: the image as witness and content credentials, in chapter 17; war-crimes evidence with fingerprints and blockchain — eyeWitness, the Starling Lab, the Berkeley Protocol — in chapter 20; the orbital witness, in chapter 104; the liar's dividend at the ballot box, in chapter 106; the symmetrical proof in the street, in chapter 109; and the chain of custody of aid — including the Venezuela earthquake — in chapters 49 and 50. Here all these pieces meet on their most extreme stage.

A siege on information, too

The first layer of the siege is the closed door. Since the start of the war, Israel has not allowed foreign journalists independent entry into Gaza; only occasional visits escorted by the army, case by case — a format the Foreign Press Association, which represents the foreign correspondents, denounces as insufficient for any real coverage [6]. The judicial chronology says it all: in January 2024, the Israeli Supreme Court rejected the FPA's first petition — there is "no vested right" to enter, it said, and the security reasons carry weight —; the FPA appealed again in September 2024; extensions followed upon extensions; and in January 2026, with the ceasefire already in force, the government told the court once more that the ban had to stay "for security reasons" — and the court postponed the decision yet again [6]. More than two dozen news organizations from around the world have jointly demanded to be allowed in [6]. The result is a historical anomaly: the first great war of this century that the world has seen almost exclusively through the eyes of those suffering it.

The second layer is what has happened to those who were inside. The CPJ counts 263 journalists and media workers killed since October 7, 2023 — plus 108 jailed and 174 wounded — and maintains that no other actor had killed so many journalists since the organization began keeping records, in 1992 [1]. The

"News Graveyards" report by Brown University's Costs of War project (April 2025) offered a perspective that is hard to forget: more journalists killed in Gaza than in the American Civil War, the two world wars, Korea, Vietnam, Yugoslavia and Afghanistan combined — a comparison to be taken with the caveat that the historical database it rests on itself advises against comparing eras counted with different methods [2]. Reporters Without Borders (RSF), which counts more than 210 dead and has reasonable grounds to believe that at least 56 were deliberately targeted for their work, has filed five complaints with the Office of the Prosecutor of the International Criminal Court — the latest on September 30, 2025 — and also documents the pattern of the double bombing: a second missile arriving once rescuers and reporters are already on the ground of the first [7].

Two names condense the mechanism. Anas al-Sharif, an Al Jazeera correspondent, was killed on August 10, 2025 in the strike on the press tent outside al-Shifa Hospital, along with four colleagues; the Israeli army declared him a "Hamas operative," an accusation Al Jazeera and the CPJ reject. The detail that matters comes earlier: back in July, the CPJ and United Nations special rapporteur Irene Khan had already warned of the smear campaign against him and of the risk to his life [4]. First the doubt, then the missile. And Wael al-Dahdouh, Al Jazeera's Gaza bureau chief, learned live on air, while reporting, that an airstrike had just killed his wife, his son Mahmoud, his daughter Sham and his grandson Adam, on October 25, 2023. He kept reporting [5].

And when the press from outside is at the door and the press inside is under the bombs, the void is filled by the witness. The war has been told to the world by figures like Bisan Owda — a twenty-five-year-old reporter with four million followers who in September 2024 won an Emmy for "It's Bisan From Gaza — and I'm Still Alive," despite a campaign to discredit her — or the photojournalist Motaz Azaiza, whose images traveled around the world [8]. In Gaza, the border between the professional photojournalist and the citizen witness has been erased: they are often the same person, with the same phone. And this makes the question of this chapter unavoidable: what evidentiary force does what they record carry?

"Fake": the liar's dividend in wartime

The book has seen the liar's dividend at the ballot box (chapter 106) and knows that detecting is not enough (chapter 15). War teaches its crudest form: there is no need to falsify anything — it is enough to say "fake" to everything. "Pallywood," the accusation that Palestinians stage their suffering, went through a measurable explosion: BBC Verify counted a peak of 220,000 uses of the term on X in November 2023 alone [9]. The "evidence" accompanying it was usually recycled material — an Egyptian protest from 2013, a Thai Halloween costume, an advertising shoot in Algeria — presented as if it were Gaza [9].

Two cases show the two faces of the mechanism. The first: the Jerusalem Post published — and later had to retract — the claim that the body of a five-month-old Palestinian baby was a doll; Israel's official account on X amplified the accusation; the BBC tracked down the family and the photographers and confirmed it was the real body of a child [9]. The second, inverse and even more corrosive: a photo released by Israel of the charred body of a baby murdered on October 7 was massively branded "AI-generated" because a commercial detector, "AI or Not," returned a positive. The company ended up acknowledging the false positive — the fault of compression and a blurred label — and several experts considered the photo genuine [9]. The lesson cuts both ways: detectors get it wrong with fake photos and with real ones, and a bad detector is the denialist's dream — any authentic image can be "debunked" with a free tool and a screenshot.

The sociologist Ron Dudai has described it precisely: the old conspiracy theories, elaborate and laborious, have given way to a cruder form of denialism — documentation is dismissed with a single word, "Fake" [9]. And the discredit does not stop at rhetoric: RSF documents how the accusations of "staging" against Gazan journalists — what it calls "Gazawood" — turn them into targets [7]. It is the path al-Sharif walked: the doubt sown in July preceded the missile of August [4].

The liar's dividend flies no flag: it is collected by any power with an interest in an image not existing — in Gaza, in Syria, anywhere — and it is always paid by the same person: the witness who tells the truth and cannot prove it. Against a "fake!" that costs zero, the answer cannot be an expert report that costs months. It must be a proof that already existed.

When the siege is at home: Minneapolis, winter 2026

It does not take a war zone or a dictatorship for the mechanism to fire. In December 2025, the United States Department of Homeland Security unleashed Operation Metro Surge on Minneapolis and Saint Paul — "the largest immigration enforcement operation ever carried out," in the department's own words —: thousands of agents from Immigration and Customs Enforcement (ICE) and the Border Patrol, more than four thousand at the peak according to figures a judge had to force into the open, and close to three thousand eight hundred arrests in two months, less than a quarter of them with criminal records [21]. The city was, precisely, the one where a teenager's phone video had changed history in 2020 (chapter 109). What happened there is the cleanest experiment in existence on this book's question: what is the witness worth when the one denying is the state itself.

Renée Good: three shots and too many cameras

On the morning of January 7, 2026, Renée Good — thirty-seven years old, a writer, a US citizen — was at the wheel of her car on a street in south Minneapolis, stopped perpendicular to the path of some federal vehicles; her partner, a six-year-old child and the dog were also aboard [22]. ICE agent Jonathan Ross fired three shots at her. The official version arrived within hours and from the very top: the secretary of Homeland Security, Kristi Noem, stated that Good had tried to "run over and ram" the agents and that she had spent the whole day "harassing and obstructing" the operation — no evidence of the latter has ever been produced — and President Trump wrote that she had struck the agent "in a violent, deliberate and vicious manner" and that he "was recovering in the hospital" [22].

The videos said otherwise, and there were many: from witnesses, from surveillance cameras and from Ross's own phone. The New York Times reconstruction fixed the sequence — one shot into the windshield and two more, at close range, through the driver's window, with the car turning to move away and the agent standing a few steps off, out of its path — and another angle showed Good, a few minutes earlier, yielding to the federal vehicles and waving them through [22]. And one detail worth an entire treatise: Ross's own recording jerks upward at exactly the decisive moment, leaving it out of frame [22]. The state's record went blind precisely when it was needed most; the citizens' records did not.

Alex Pretti: the witness who was shot

Seventeen days later — the day after a general strike had filled the city center [21] — the operation killed Alex Pretti: thirty-seven years old, an intensive-care nurse at the veterans hospital, a US citizen, with no record and a gun permit in order, as the Minneapolis chief of police confirmed [23]. That morning Pretti was one of the people documenting an ICE operation with his phone — the exact figure of the witness from chapter 109. Secretary Noem declared that he had approached the agents "with a gun"; the department added that he wanted to "do maximum harm and massacre law enforcement" and that he carried "two magazines and no identification of any kind"; the White House deputy chief of staff, Stephen Miller, called him a "domestic terrorist" [23].

The videos of those present — multiple angles, reviewed and verified by Reuters, the BBC, the Wall Street Journal, AP, CNN, the New York Times and The Guardian — show a different sequence: Pretti with his phone in his right hand; pepper-sprayed while trying to help a woman who had fallen; half a dozen agents pinning him to the ground; one agent pulling the pistol from the holster on his hip — and, a second later, another firing some ten shots at point-blank range in five seconds, the last ones when he already lay motionless [23]. No agent attempted resuscitation, and a pediatrician who offered was pushed away [23]. The county medical examiner ruled the death a homicide [23]. And the names of the two agents were not made public by the government: ProPublica published them [24].

The siege on the evidence

Density did its work. The official versions did not last even a week: the mayor and the governor, after seeing the videos, publicly contradicted the federal account [23]; one federal judge counted at least ninety-six court orders violated by ICE in a month, and another found that the "overwhelming majority" of those arrested were in the country legally [21]; and on February 12, the White House "border czar" announced the end of the operation at a press conference where he avoided speaking the two names [21].

But the same machinery closed ranks around the proof: the day after the first shooting, the FBI revoked the state criminal investigation agency's access to the evidence [22]; by April, Pretti's phone was still being held, Good's car remained sealed in a federal warehouse without a single state investigator having been able to examine it, and the agents' names and files were being systematically denied [24]. Minnesota and Hennepin County had to file suit to obtain evidence that — they maintain — is being withheld "to shield

the agents involved" [24]. As we write this, no agent has been charged with anything: the man who shot Renée Good spent three days off duty and months later was operating in another state [22], and the only known disciplinary measure in the entire operation is two agents on administrative leave for "false statements" about the third person shot, a man wounded by a bullet in the leg [24]. The county attorney keeps seventeen criminal investigations open against a wall of immunities [24]. "It looks like they are actively preventing the state from investigating," was how one law professor summed it up [24].

The pattern: Chicago, Houston

Minneapolis is not an anomaly: it is the best-documented case of a pattern. Three months earlier, in Chicago, an agent had fired five shots at Marimar Martinez, a teacher who was following the federal vehicles honking her horn to warn the neighbors — and the one accused of ramming was her; the body camera ended up showing the agents' vehicle turning into hers right after one of them said "it's time to get aggressive," along with the agent's own messages boasting about it — "five shots and seven holes" —; the charges against her were dismissed with prejudice [25].

That same November, federal judge Sara Ellis confronted, in a two-hundred-and-twenty-three-page opinion, the official use-of-force reports with the body-camera videos she herself had ordered activated: the government's account was "simply not credible," the head of the operation had kept changing the order of events under oath — the gas fired before the stone he claimed to have received — and one agent had drafted the official report of an incident with ChatGPT from "a brief sentence about an encounter and several images" [25]. It is worth pausing on this last detail: in the dispute over reality between citizen and state, the synthetic account was the official one. The running count since January 2025 reaches thirty-eight shootings by immigration agents, with nine dead [25]. And it does not stop: on July 7, 2026, as this chapter was being finished, a fifty-two-year-old man died in Houston from a shot fired by an ICE agent who alleges — once again — a ramming; the agents wore no body cameras, the accounts of the preceding moments contradict one another, and the county medical examiner has ruled the death a homicide [26].

The counterfactual: a single witness, no anchor

Now, the exercise this book cannot spare itself. Replay January 7 with a single variation: there is no crowd. No surveillance camera pointed at the street, no neighborhood gathered; only someone who, frightened, records it in secret from a window — and takes weeks to make it public, out of fear. By then, the official version has already been "the truth" for a month: it stands in statements, in hearings, in the president's mouth. When the video comes out, the manual this chapter has already described will fire on its own: "staged," "actors," "it's AI" — and a free detector and a false positive like the one we have already seen [9] will be enough to manufacture the "debunking."

Against the sworn word of a state, a file without a certain date or provable integrity is just that: a file. Because what saved the truth in Minneapolis was no institution: it was density — too many cameras, too many witnesses, too many angles to deny them all at once. And density is an accident: it was there because the operation was massive and contested, and the community was forewarned and watching. The next Renée Good may have a single witness; the next Alex Pretti, none. And where no counterweights remain — no judge counting violated orders, no county attorney able to sue, no general strike — the official version does not so much win as become the only one in existence. In a state like that, openly dictatorial, the witness at the window does not have a video: they have an accusation of forgery waiting for them.

With proof of existence, the same scenario changes physics. If the clandestine video was sealed at birth — the fingerprint, and only the fingerprint, left the phone for the public anchor, dated that same January 7 [20] — the witness can wait weeks or years, or reveal it from exile: their anchor will always predate the official account, and "this was fabricated afterward" stops being an exit. That is exactly what no registry dependent on power can offer: the agent's camera that jerks upward at the decisive moment, the phone held back, the car sealed away — every convenient gap in the official record becomes visible when the rest of the evidence carries a certain date and verifiable integrity. And the symmetry this chapter claims further on holds here in both directions: the same seal would have protected an agent who told the truth. It asks no one to be trusted; it only makes lying expensive, in every direction.

When the crowd is there, truth can win by density; when there is only one witness, everything rides on whether their proof was born sealed. And when there is neither crowd nor seal, what remains is the slow craft of rebuilding the facts after the fact — a heroic craft whose price, in time and in doubt, is what comes next.

The forensics of afterward: it wins, but late

Independent forensic reconstruction is the heroine of this information war. And at the same time, the demonstration of how much arriving late costs. Three cases, in gradation.

Shireen Abu Akleh: four months to topple a version

Shireen Abu Akleh, an Al Jazeera correspondent, died from a shot to the head on May 11, 2022 in Jenin, wearing a helmet and a PRESS vest. The first official version pointed to indiscriminate Palestinian fire. What came next is the exact cadence of the rebuttal: CNN's investigation (May 24) found no active combat nearby and, with witnesses and audio analysis, pointed to Israeli fire; AP, the New York Times and the Washington Post published matching reconstructions; the UN Human Rights Office concluded in June that the shots came from Israeli forces — "several single, seemingly well-aimed shots" —; the United States assessed in July that Israeli fire was most likely, with a bullet too damaged for conclusive ballistics; and on September 5, 2022 the Israeli army admitted that one of its soldiers had "probably" killed her. No charges [11]. Almost four months and half a dozen independent investigations to topple the first version — with access to the ground, in relative calm and with a famous victim. In Gaza, without access, this path multiplies — or becomes impossible.

Hind Rajab: 355 impacts and a phone call

Hind Rajab was six years old. On January 29, 2024 she was trapped in a car under fire in Gaza City, surrounded by the bodies of her family, and begged for help by phone from the Palestine Red Crescent; her body, those of six relatives and those of the two paramedics from the ambulance that had gone out to find her appeared two weeks later. The reconstruction by Forensic Architecture with the NGO Earshot and Al Jazeera established that the car had taken 355 bullet impacts; and the acoustic analysis of the call — the interval between the sound of each impact and that of the muzzle blast — placed the shooter at between 13 and 23 meters [12]. It is worth pausing on what the central piece of evidence is here: the audio of a mobile phone call. The spontaneous digital witness already decides international investigations; what it still lacks is a way to be born indisputable.

The Caesar photos: a decade to Koblenz

And the precedent that closes the circle comes from Syria, from a dictatorship: a forensic photographer of the Syrian military police, code name "Caesar," fled in 2013 with some 53,000 photographs of people who had died under torture in the regime's prisons. Authenticating them took years — expert analyses of the metadata of the regime's own system, which numbered devices and records, and a forensic medical analysis of thousands of images by experts at the University of Cologne, which found systematic torture and death [13]. On January 13, 2022, the court in Koblenz, Germany, trying the case under universal jurisdiction, sentenced Anwar Raslan, former head of interrogations of Branch 251, to life imprisonment for crimes against humanity — with the Caesar photos as a key piece of evidence [13]. Between the mortal risk of one man carrying hard drives out of a country at war and the verdict: almost a decade. Courage saved the proof; the system took ten years to be able to trust it.

Courts that wait, evidence that never arrives

This evidence, when it arrives, already has somewhere to go. The International Criminal Court issued arrest warrants on November 21, 2024 against Benjamin Netanyahu and Yoav Gallant — for the war crime of starvation as a method of warfare and for crimes against humanity — and against the Hamas commander Mohammed Deif [14]; the International Court of Justice, in the case brought by South Africa under the Genocide Convention, ordered binding provisional measures in January, March and May 2024 [15]; and the framework for accepting open-source digital evidence — the Berkeley Protocol — already exists (chapter 20). There is no shortage of courts. What is missing is for the evidence to arrive whole, fast and indisputable.

Because when it does not arrive, what happened with al-Ahli Hospital happens. The explosion of October 17, 2023 — 471 dead according to the Gaza Ministry of Health; between 100 and 300 according to US intelligence — set the best forensic teams in the world working on the same videos. And they diverged: AP, CNN and other outlets concluded that a failed rocket launched from Gaza was most likely; the New York Times ended up saying the cause was uncertain; Human Rights Watch saw the pattern of a propelled rocket but demanded an independent investigation; Forensic Architecture concluded that the projectile came from the direction of Israel [10]. Without access to the crater, with recompressed and cropped

videos, without a certain time or verifiable integrity for the key files, the same evidence supported opposite conclusions — and each side kept its own. Forensics did not fail: it lacked the foundations. Exactly the ones a seal at the origin would have given it.

Sealing at the origin: the accredited press under siege

That a newsroom can seal what it publishes is no longer a hypothesis. Reuters and the Starling Lab tried it in 2020 and 2021 with "78 Days": over the seventy-eight days of the US presidential transition — storming of the Capitol included — more than two hundred photographs were captured with cryptographic metadata and anchored on public chains, with the provenance of every image verifiable by any reader [18]. If it could be done for a political transition in Washington, it can be done for a war. That is exactly where it is needed most.

The t2p flow for the press begins on the device of whoever presses the shutter: the SHA-512 fingerprint — computed on the phone, from the exact bytes of the file — receives its RFC 3161/eIDAS timestamp and is anchored on the public blockchain; the image does not leave the device to generate the proof [20]. The GPS position is added when it is reliable. And all of it travels in an .asice package verifiable offline, from the conflict zone to the newsroom and, if need be, to the prosecutor's office — the same chain of custody that chapter 50 applies to aid, applied to the witness [20]. It is, in fact, what the Berkeley Protocol will demand of any piece of evidence tomorrow (chapter 20), generated by default today.

The credential solves the other half of the problem: the agency signs the link — "this material is from an accredited contributor of ours" — without publishing who they are or where they are, the selective disclosure of chapter 20 [20]. The newsroom can prove provenance to the reader and to the judge, and protect the person at the same time. And if they are killed — in Gaza this is not a figure of speech [1] — their sealed archive goes on proving what they saw and when. Killing the messenger stops killing the message.

And here is where this book's doctrine becomes literal. Content credentials signed in the camera (chapter 17) are the perfect complement — but they are metadata that travel with the file, and a recompression wipes them; and every certification system has an owner, and every owner has a door where someone can knock. The anchor on a public chain has none: no government — including the one denying the images — can unpublish a fingerprint or backdate it. Chapter 112 asked who controls the proof; here the answer decides lives: when the accused is a state, the proof cannot depend on any authority that state can pressure. The chain's neutrality is not an abstract virtue — it is the technical requirement for the proof to hold against any power, one's own included.

The witness with the phone: the same proof for everyone

The idea that the citizen witness deserves tools of proof is not new. WITNESS and the Guardian Project launched ProofMode in 2017: a free app that signs and fingerprints every photo and video at the moment of capture, with a snapshot of the sensors, and notarizes its fingerprint — with OpenTimestamps, the Bitcoin anchoring the book already knows, among its notary systems [19]. And eyeWitness to Atrocities has done something similar since 2015 for trained documenters, with more than 20,000 items already used as evidence in court (chapter 20). What was missing is not the idea: it is for it to be universal, by default, and designed for someone who is neither an activist nor a journalist — for anyone who one morning finds themselves recording the worst day of their life.

For this scenario — the worst — the seal must have four properties, and none of them is optional.

Only the fingerprint leaves the phone. The photo is not uploaded anywhere to generate the proof: the device computes the SHA-512 fingerprint and it is this — this alone — that receives the timestamp and the anchor on the public blockchain [20]. At a military checkpoint, in a search, there is no upload to give the witness away; the image can live encrypted or in a copy off the device, and the anchored fingerprint bides its time.

Seal now, publish when it is safe. This is the commit-and-reveal of chapters 45 and 46, applied to survival: the victim seals the video's fingerprint today and reveals it in a month, a year, or from exile — and the certain date that counts is today's. The denialist who says "this was fabricated afterward" collides with an anchor that predates their own denial.

The posthumous proof. If the witness dies — and 263 times is not a hypothesis [1] — the seal survives: anyone holding the copy can prove what existed and when. The archive of every Bisan and every Motaz [8] becomes indestructible by decapitation: you can kill the one who records; you cannot kill the fact that it was already recorded.

Symmetry. This proof takes no sides: the video of an Israeli civilian on October 7 and that of a family in Khan Younis under the bombs carry exactly the same anchoring force. It is the symmetrical proof of chapter 109: Darnella Frazier's camera counted because it was authentic, not because of whose it was. The technology that shelters the Palestinian witness shelters the Israeli, the Ukrainian and the Sudanese alike — and for that reason, and that reason alone, it can aspire to be a global standard.

The ethical limit remains, and it is non-negotiable (chapters 49 and 50): never mandatory biometrics for victims. Proof of life and the video selfie are options for those who want to claim authorship — not a toll. Whoever documents a crime must not pay with their identity; and geolocation is revealed when the witness decides, not before (chapter 20).

The promise and the peril

The promise is a war where the access ban no longer manufactures darkness, because every phone inside is a notary [19][20]; where the reflexive "fake!" is answered with a dated fingerprint and not with months of expert reports [9]; where the next al-Ahli is settled on videos with a certain time and verifiable integrity [10]; where the dead journalist's archive goes on accusing [1]; and where the ICC prosecutor's office receives evidence packages that meet the Berkeley Protocol by default [14] (chapter 20). The truth will still cost lives; at least let it not be buriable with them.

The possession that kills — a phone with a sealed file can be a death sentence at a search. The design must assume it: only fingerprints off the device, content encrypted or off the phone, deferred revelation [20]. But no architecture removes the witness's physical risk — and promising invulnerability would be lying. The seal protects the proof; protecting the person remains the work of the law, of the media that publish — and the responsibility of whoever wants them silenced.

The blackout — without a network there is no anchor: Gaza has accumulated at least twenty-three documented communications blackouts since October 2023, from strikes on the fiber and from lack of fuel — civilian connectivity turned into a means and method of warfare, in the words of Access Now [16]. The local seal is generated all the same and anchored when the signal returns; but then the chain only proves the "not after." The partial remedy is the temporal sandwich: folding into the capture a recent, unpredictable public value — the fingerprint of the last block known before the blackout — proves the "not before," and the later anchor, the "not after." The window is closed on both sides; and an honest bracket of days is worth more than a false promise of instantaneity.

The GPS that lies — in the region, jamming and spoofing of the GPS signal are massive and documented: more than 50,000 flights affected in 2024, civilian apps placing users in Beirut or Cairo [17]. An honest system acknowledges it: when the signal is not reliable, the position is degraded or omitted — and the seal's core value, certain time and integrity, does not depend on it [20]. A spoofed geolocation presented as certain would do more harm than no geolocation at all.

The seal does not sanctify — proof of existence attests when a file existed and that no one has touched it: it does not attest what was happening outside the frame, nor does it prevent sealing a staged scene the same day. Whoever wants to stage theater with a certain date will be able to. The answer is the same as always in this book: corroboration — more witnesses, more angles, the satellite (chapter 104), forensics (chapter 15) — on a base that at least is no longer in dispute: which files, from when, untouched. The seal replaces neither journalism nor expert analysis: it spares them the sterile half of the work.

The registry that hunts — the same trail that shelters the witness can de-anonymize them if it is badly designed (chapter 20): anchors linkable to one another, metadata that draws patterns, geolocation by default. Selective disclosure, non-correlatable identifiers and the witness's control over every field are not luxury extras: they are the difference between a human rights tool and a detection list at the persecutor's disposal.

It's time to act

The siege on truth is fought before it begins — when the photo is born.

For agencies and news media: the "78 Days" precedent has already shown you the way [18] — make sealing at capture the standard for all conflict-zone material, your own or a contributor's; publish verifiable provenance with every piece; and sign your local contributors' credential without exposing them [20]. Their risk is already the maximum: let their word, at least, be impossible to erase.

For photojournalists and contributors inside the zone: seal everything, even what you do not publish — commit-and-reveal [20] lets you document today and reveal when you are out; and your agency or your family will be able to prove it even if you are gone. The PRESS vest does not stop missiles; the anchored fingerprint does not stop lies — but it makes them expire.

For human rights organizations and international prosecutors: demand and teach sealing at the origin — investigations like those of Abu Akleh, Hind Rajab or the Caesar files [11][12][13] would have gained months or years; the Berkeley Protocol (chapter 20) already gives you the framework for accepting it; and every workshop on safe documentation in a conflict zone is worth more than ten statements after the fact.

For citizens and victims: your phone is a notary that asks no one's permission [19][20]. Seal what you see — only the fingerprint leaves the device; publish when it is safe; and if you choose anonymity, the proof will count just the same. No one may demand your identity as the price of the right to be believed.

For the platforms: preserving the credentials and fingerprints of conflict material instead of wiping them on recompression (chapter 17) is the difference between being the witness's channel or their shredder; and labeling verified provenance — without censoring what lacks it — gives the reader the criterion detectors cannot give them [9].

For democratic governments and international bodies: if the commitment to a free press is real, it holds even when the war is yours or an ally's — access for the accredited press [6], protected civilian connectivity [16], and procedural recognition of proof of existence, so that the witness of the next war does not depend on the courage of one official fleeing with 53,000 photos [13].

In the end, this whole book — and this chapter more than any other — defends something old: that one person's testimony counts, and that their word deserves to be believable. Technology only adds what force has always tried to take from it: the proof. For all those who kept recording until the last second, time2act.

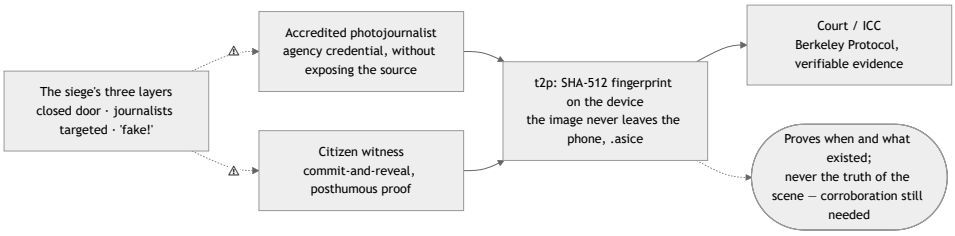


Figure 113.1 — The proof under siege, from the lens to the courtroom (t2p: SHA-512 fingerprint on the device — the image never leaves the phone —, RFC 3161/eIDAS timestamp, anchor on the public blockchain, .asice package): the accredited photojournalist with an agency credential that does not expose the source; the citizen witness with commit-and-reveal (seal now, publish when it is safe) and posthumous proof; the three layers of the siege (door closed to the international press, local journalists killed or marked, "fake!" against real images) and where the seal breaks them. Limits: the blackout delays the anchor (the "not before / not after" temporal sandwich); spoofed GPS degrades the position, not the time; the seal proves when and what — never the truth of the scene: corroboration remains indispensable. Sources: CPJ; RSF; Costs of War (Brown); OHCHR; Forensic Architecture; ECCHR/Koblenz; Access Now; Starling Lab/Reuters; WITNESS/Guardian Project.

---

[1] CPJ (Committee to Protect Journalists) — *Israel-Gaza War* (263 journalists and media workers killed since 7-Oct-2023; 108 jailed; 174 wounded; "the deadliest and most deliberate effort to kill and silence journalists" the CPJ has ever documented) and *At least 5 Gazan journalists killed in Israeli strike on Nasser Hospital* (25-Aug-2025). https://cpj.org/issue/israel-gaza-war/ · https://cpj.org/2025/08/at-least-5-gazan-journalists-killed-in-israeli-strike-on-nasser-hospital/ ↩

[2] Watson Institute (Brown University), Costs of War project — *News Graveyards: How Dangers to War Reporters Endanger the World* (Nick Turse, 1-Apr-2025: more journalists killed in Gaza than in the great wars of the twentieth century combined — a comparison the underlying database itself qualifies; ~13 Palestinian journalists killed per month); Foreign Policy — *War in Gaza the 'Worst Ever Conflict for Reporters'*. https://watson.brown.edu/costsofwar/papers/2025/Journalists · https://foreignpolicy.com/2025/04/01/gaza-war-worst-ever-for-reporters-costs-of-war-project/ ↩

[3] Wikipedia (EN) — *2025 Nasser Hospital strikes* (25-Aug-2025; first strike ~10:00 against the position of the Reuters camera — which was broadcasting live — and second strike ~10 min later against rescuers and journalists; ≥22 dead, 5 journalists: Hussam al-Masri/Reuters, Mariam Dagga/AP, Mohammed Salama/Al Jazeera, Moaz Abu Taha, Ahmed Abu Aziz; "tragic

mishap"); CNN — *New video reveals third strike* (27-Aug-2025). https://en.wikipedia.org/wiki/2025_Nasser_Hospital_strikes · https://www.cnn.com/2025/08/27/middleeast/gaza-nasser-hospital-israel-attack-three-strikes-intl ↩

[4] CPJ — *Israel kills Al Jazeera journalists in targeted Gaza City airstrike* (Anas al-Sharif and four colleagues, 10-Aug-2025; the prior alert by the CPJ and UN rapporteur Irene Khan, Jul-2025); NPR — *Israeli airstrike kills Al Jazeera's Anas al-Sharif in Gaza*. https://cpj.org/2025/08/israel-kills-al-jazeera-journalists-in-targeted-gaza-city-airstrike/ · https://www.npr.org/2025/08/12/g-s1-82143/gaza-anas-al-sharif-al-jazeera-israel ↩

[5] Al Jazeera — *Family of Al Jazeera Gaza bureau head killed in Israeli air raid* (25-Oct-2023: Wael al-Dahdouh's wife, son Mahmoud, daughter Sham and grandson Adam) and *To kill a family: the loss of Wael Dahdouh's family*. https://www.aljazeera.com/news/2023/10/25/family-of-al-jazeera-gaza-bureau-head-killed-in-israeli-air-raid · https://www.aljazeera.com/news/2023/11/1/to-kill-a-family-the-loss-of-wael-dahdouhs-family-to-israeli-bombs ↩

[6] The Times of Israel / Al Jazeera / Courthouse News — the ban on international press access to Gaza: the FPA's petition rejected (Jan-2024, "no vested right"), new petition (Sep-2024), the government keeps the ban "for security reasons" despite the ceasefire (Jan-2026) and the Supreme Court postpones again (28-Jan-2026); more than two dozen news organizations demand access. https://www.timesofisrael.com/ban-on-media-access-to-gaza-should-stay-in-place-israel-tells-high-court/ · https://www.aljazeera.com/news/2026/1/7/media-body-condemns-israels-continued-ban-on-foreign-media-access-to-gaza · https://www.aljazeera.com/news/2026/1/28/israels-top-court-delays-gaza-press-access-ruling-amid-years-long-ban · https://www.courthousenews.com/over-two-dozen-global-news-organizations-call-on-israel-to-allow-gaza-access/ ↩

[7] RSF (Reporters Without Borders) — *RSF files fifth complaint with ICC about Israeli war crimes against journalists in Gaza* (30-Sep-2025; five complaints since Oct-2023; ≥210 journalists killed, 56 with reasonable grounds to believe they were deliberately targeted; the "double bombing" pattern) and *"Gazawood": the deadly accusations against Gaza's journalists*. https://rsf.org/en/rsf-files-fifth-complaint-icc-about-israeli-war-crimes-against-journalists-gaza · https://rsf.org/en/gazawood-deadly-accusations-against-gaza-s-journalists ↩

[8] Al Jazeera / The National — *Palestinian journalist Bisan Owda and AJ+ win Emmy for Gaza war documentary* (25-Sep-2024, "It's Bisan From Gaza — and I'm Still Alive"; ~4 million followers; a failed discrediting campaign) and the photojournalist Motaz Azaiza. https://www.aljazeera.com/news/2024/9/26/palestinian-journalist-bisan-owda-and-aj-win-emmy-for-gaza-war-documentary · https://www.thenationalnews.com/arts-culture/film-tv/2024/09/26/bisan-owda-emmy-award-palestine/ ↩

[9] Wikipedia (EN) / NPR — *Misinformation in the Gaza war* and *Pallywood* (a peak of 220,000 uses on X in Nov-2023 according to BBC Verify; the "doll" case published and retracted by the Jerusalem Post and amplified by Israel's official account on X; the false positive by the "AI or Not" detector on a real photo; Ron Dudai: documentation dismissed with a single word, "Fake"); NPR — *False 'crisis actor' claims spread about Israel-Hamas war*. https://en.wikipedia.org/wiki/Misinformation_in_the_Gaza_war · https://en.wikipedia.org/wiki/Pallywood · https://www.npr.org/2023/11/27/1214451419/civilian-deaths-are-being-dismissed-as-crisis-actors-in-gaza-and-israel ↩

[10] Human Rights Watch / CNN / Forensic Architecture — the explosion at al-Ahli Hospital (17-Oct-2023; 471 dead according to the Gaza Ministry of Health, 100-300 according to US intelligence) and the forensic divergence: HRW (pattern consistent with a propelled rocket; an independent investigation is needed), CNN and AP (a failed rocket from Gaza, most likely), NYT (cause uncertain), Forensic Architecture (projectile from the direction of Israel). https://www.hrw.org/news/2023/11/26/gaza-findings-october-17-al-ahli-hospital-explosion · https://www.cnn.com/2023/10/21/middleeast/cnn-investigates-forensic-analysis-gaza-hospital-blast · https://forensic-architecture.org/investigation/israeli-disinformation-al-ahli-hospital ↩

[11] OHCHR / CNN — the killing of Shireen Abu Akleh (11-May-2022): the UN's conclusion (Jun-2022: shots from Israeli forces, "seemingly well-aimed"), CNN's investigation (24-May-2022) and the Israeli army's admission (5-Sep-2022: "probably" its own fire; no charges). https://www.ohchr.org/en/press-briefing-notes/2022/06/killing-journalist-occupied-palestinian-territory · https://www.cnn.com/2022/05/24/middleeast/shireen-abu-akleh-jenin-killing-investigation-cmd-intl · https://www.cnn.com/2022/09/05/middleeast/idf-shireen-abu-akleh-investigation-intl/index.html ↩

[12] Forensic Architecture / Earshot — *The Killing of Hind Rajab* (29-Jan-2024: 355 impacts on the car; acoustic analysis of the call: shots from 13-23 m). https://forensic-architecture.org/investigation/the-killing-of-hind-rajab · https://www.earshot.ngo/investigations/the-killing-of-layan-hamada-and-hind-rajab ↩

[13] ECCHR / Syria Justice & Accountability Centre / CBS — the Caesar files (~53,000 photos carried out of Syria in 2013; expert authentication of metadata and medical forensics) and the Koblenz verdict (13-Jan-2022: life imprisonment for Anwar Raslan for crimes against humanity). https://www.ecchr.eu/en/case/caesar-photos-document-systematic-torture/ · https://syriaaccountability.org/inside-the-raslan-trial-a-forensic-analysis-of-the-caesar-photos/ · https://www.cbsnews.com/news/syria-war-bashar-assad-police-officer-crimes-against-humanity-germany/ ↩

[14] International Criminal Court — *ICC Pre-Trial Chamber I rejects the State of Israel's challenges to jurisdiction and issues warrants of arrest for Benjamin Netanyahu and Yoav Gallant* (21-Nov-2024; the war crime of starvation as a method of warfare; crimes against humanity; a warrant also for Mohammed Deif). https://www.icc-cpi.int/news/situation-state-palestine-icc-pre-trial-chamber-i-rejects-state-israels-challenges ↩

[15] International Court of Justice — *Application of the Genocide Convention (South Africa v. Israel)*, provisional measures of 26-Jan-2024, 28-Mar-2024 and 24-May-2024 (binding). https://www.icj-cij.org/case/192 ↩

[16] Access Now (#KeepItOn) / EFF — *Israel must stop cutting Gaza off from the world* (19-Jun-2025: ≥23 communications blackouts since Oct-2023, from strikes on the fiber and lack of fuel; "civilian connectivity turned into a means and method of warfare") and *Connectivity is a Lifeline, Not a Luxury*. https://www.accessnow.org/press-release/keepiton-israel-must-stop-cutting-gaza-off/ · https://www.eff.org/deeplinks/2025/06/connectivity-lifeline-not-luxury-telecom-blackouts-gaza-threaten-lives-and-digital ↩

[17] NPR / GPS World / The Times of Israel — GPS jamming and spoofing in the region (2024: >50,000 flights affected; a ×6 increase; apps placing users in Beirut or Cairo; an Israeli air base identified as a source).

https://www.npr.org/2024/04/22/1245847903/israel-gps-spoofing · https://www.gpsworld.com/israeli-air-base-identified-as-alleged-source-of-gps-disruptions-in-mideast/ · https://www.timesofisrael.com/gps-jamming-amid-wars-playing-havoc-with-airline-navigation-report/ ↩

[18] Starling Lab / Content Authenticity Initiative — *78 Days* (Reuters + Starling, Nov-2020→Jan-2021: >200 photos of the presidential transition with cryptographic metadata anchored on Filecoin and Hedera; provenance verifiable by anyone). https://www.starlinglab.org/78days/ · https://contentauthenticity.org/blog/cai-spotlight-78-days-from-stanford-starling-lab-and-reuters ↩

[19] Guardian Project / WITNESS — *ProofMode* (since 2017: OpenPGP signature and fingerprint at capture, sensor snapshot, notarization — with OpenTimestamps among the supported systems; designed for citizen human rights documentation). https://guardianproject.info/apps/org.witness.proofmode/ · https://blog.witness.org/2017/04/proofmode-helping-prove-human-rights-abuses-world/ ↩

[20] t2p — product manual (SHA-512 fingerprint computed on the device — only the fingerprint leaves the phone —; RFC 3161/eIDAS timestamp; anchoring on the public blockchain; forensic GPS with honest degradation; offline mode with deferred anchoring; commit-and-reveal; .asice package verifiable offline; certificate with QR). See chapters 13 and 15; commit-and-reveal, in chapters 45 and 46; the aid chain of custody, in chapter 50; selective disclosure and war-crimes evidence, in chapter 20; content credentials, in chapter 17. ↩

[21] Wikipedia (EN) — *Operation Metro Surge* (4-Dec-2025 → Feb-2026: "the largest immigration enforcement operation ever carried out" according to the DHS; thousands of ICE and CBP agents in Minneapolis–Saint Paul; ≥3,789 arrests, under 25% with criminal records; the general strike of 23-Jan-2026; Judge Schiltz: ≥96 court orders violated by ICE since 1-Jan; Judge Blackwell: the "overwhelming majority" of those arrested had legal presence); NPR — *Minnesota ICE crackdown will end* (12-Feb-2026: Tom Homan announces the end of the operation — which "has cost the lives of two American citizens" — without naming them); Sahan Journal — the deployment figures (more than 4,000 agents at the peak), revealed under court order. https://en.wikipedia.org/wiki/Operation_Metro_Surge · https://www.npr.org/2026/02/12/nx-s1-5712280/minnesota-ice-surge-ends · https://sahanjournal.com/immigration/hundreds-ice-agents-leaving-minnesota-operation-metro-surge/ ↩

[22] Wikipedia (EN) — *Killing of Renée Good* (7-Jan-2026, Portland Avenue, Minneapolis: ICE agent Jonathan Ross fires three shots — windshield and driver's window — with the vehicle turning to move away, according to the New York Times reconstruction; in the car, her partner and a six-year-old child; the allegations by Noem and Trump; Ross's own recording, jerked upward at the critical moment; the state investigation agency's access to the evidence, revoked the next day; Ross, three days off duty and reassigned out of state); The Intercept — *Video Shows ICE Agent's Fatal Shooting of Civilian in Minneapolis* (the close-range shot at the departing vehicle; a second angle shows Good yielding to the federal vehicles minutes earlier). https://en.wikipedia.org/wiki/Killing_of_Ren%C3%A9e_Good · https://theintercept.com/2026/01/07/video-ice-shooting-civilian-minneapolis/ ↩

[23] Wikipedia (EN) — *Killing of Alex Pretti* (24-Jan-2026, Minneapolis: an ICU nurse at the veterans hospital, 37, gun permit and no record; the witnesses' videos, reviewed and verified by Reuters, BBC, Wall Street Journal, AP, CNN, New York Times and The Guardian: the phone in his hand, the pepper spray, the agents pinning him down, the gun pulled from his holster and the shots a second later; the DHS allegations — "maximum harm," "two magazines and no identification of any kind" — and Stephen Miller's; the Hennepin County medical examiner: homicide; the agents, identified by ProPublica); NPR — *Videos refute DHS account of fatal shooting* (25-Jan-2026: Noem — "approached with a gun" — against the videos: he never drew any weapon; he was filming with his phone); The Marshall Project — *How Many People Have Been Killed by ICE and Border Patrol?* (26-Jan-2026: no resuscitation by the agents; a pediatrician pushed away; in the Good case, three minutes without calling emergency services; the public rebuttal by Mayor Frey and Governor Walz after seeing the videos). https://en.wikipedia.org/wiki/Killing_of_Alex_Pretti · https://www.npr.org/2026/01/25/nx-s1-5687384/minneapolis-minnesota-shooting-video-dhs · https://www.themarshallproject.org/2026/01/26/ice-minneapolis-shooting-alex-pretti ↩

[24] ProPublica — *Minnesota's Fight to Hold Agents Who Shot Alex Pretti, Renee Good Accountable* (the lawsuit by Minnesota, Hennepin County and the state investigation agency against the Departments of Justice and Homeland Security, 24-Mar-2026: evidence withheld "to shield the agents involved" — names, weapons, reports, GPS data and witness statements denied despite formal demands; County Attorney Moriarty's seventeen criminal investigations; the supremacy-clause obstacle; the agents in the Pretti case, identified by ProPublica on 1-Feb); NPR — *Months after the ICE shootings in Minnesota, a federal probe remains elusive* (10-Apr-2026: Pretti's phone, still held; Good's car, sealed in an FBI warehouse unexamined; state investigators physically kept from the scene; the operation's only disciplinary measure: two agents on administrative leave for "false statements" in the case of the third person shot, Julio Sosa-Celis; "It looks like they are actively preventing the state from investigating" — Rachel Moran, law professor). https://www.propublica.org/article/minnesota-trump-ice-shooting-lawsuit-alex-pretti-renee-good · https://www.npr.org/2026/04/10/nx-s1-5775847/alex-pretti-renee-good-ice-shootings-federal-investigations ↩

[25] Wikipedia (EN) — *Shooting of Marimar Martinez* (Chicago, 4-Oct-2025: five shots by Border Patrol agent Charles Exum; the body camera shows the agents' vehicle turning into hers after "it's time to get aggressive"; the agent's messages — "fired 5 shots and she had 7 holes" —; the charges against Martinez and her co-defendant dismissed with prejudice on 20-Nov-2025) and *List of shootings by U.S. immigration agents in the second Trump administration* (the running count: 38 shootings and 9 dead between 20-Jan-2025 and 7-Jul-2026; the cases with charges against civilians dropped when video contradicted the agents' version); AP (via KSAT) — *Judge's blistering opinion details use of force in Chicago-area immigration crackdown* (Nov-2025: Judge Sara Ellis's 223-page opinion — the government's account, "simply not credible," contradicted by the body cameras she herself had ordered activated; Bovino's shifting versions about the stone and the gas; the use-of-force report drafted with ChatGPT). https://en.wikipedia.org/wiki/Shooting_of_Marimar_Martinez · https://en.wikipedia.org/wiki/List_of_shootings_by_U.S._immigration_agents_in_the_second_Trump_administration · https://www.ksat.com/news/politics/2025/11/21/judges-blistering-opinion-details-use-of-force-in-chicago-area-immigration-crackdown/ ↩

[26] KPRC 2 / Click2Houston — *ICE agent shoots, kills Lorenzo Salgado Araujo during attempted arrest in east Houston* (7-Jul-2026, Magnolia Park: ICE alleges he "rammed an official vehicle" and that the agent fired "in self-defense"; the agents wore no body cameras, according to the DHS; conflicting accounts of the preceding moments; the Harris County medical examiner, 9-Jul: homicide). https://www.click2houston.com/news/local/2026/07/07/heavy-law-enforcement-presence-reported-in-east-houston-neighborhood/ ↩

The Dignity of Those Who Exist: "Magnifica Humanitas" and the Proof of Existence

"The person ends up reduced to a means for obtaining results, a resource to be used and exploited."

— Leo XIV, encyclical Magnifica Humanitas, §51 (2026)

On May 25, 2026, at the Vatican, something happened that no protocol had foreseen: for the first time in history, a pope personally presented an encyclical to the world. And Leo XIV did not appear alone. Beside him sat **Chris Olah**, cofounder of Anthropic — one of the leading companies in artificial intelligence — and a **declared atheist**: raised in Toronto in a devout evangelical family, he had stopped believing at fifteen [8]. The text, signed ten days earlier and titled "Magnifica Humanitas," was presented to the world with an image as unusual as it was deliberate: a believer and a non-believer, side by side, saying at the same time that every person is unique and irreplaceable [4][6].

The pope summed it up with a sentence that speaks for this whole book: "what a great sign of hope that, with our differences, we can listen to one another" [6]. Let us begin, then, here, and not with an argument from faith. The strength of the idea that runs through the encyclical — "the dignity that belongs to every human being by the simple fact of existing" (§52) [1] — lies not in who pronounces it, but in the fact that it can be subscribed to equally by a pontiff and by an engineer who does not believe in God. The dignity of those who exist demands no creed: it precedes any belief and, precisely for that reason, can be the foundation of a technology designed for everyone.

The encyclical is published exactly 135 years after Leo XIII's "Rerum Novarum," the founding text of Catholic social doctrine that laid the bases of respect for the dignity of the worker [2]. This chronological coincidence underlines a historical continuity: just as the old document denounced the conditions of industrial exploitation, the new encyclical addresses a world dominated by algorithms and data. Its central theme — "the guardianship of the human person in the age of artificial intelligence" — echoes the call the Church itself made more than a century ago, now transposed to the digital context [1].

At the heart of the text, in paragraph fifty-two, appears a sentence that reads like a keystone for the whole debate we have conducted from chapter 10 to the chapter of the seal (ch. 112): "the dignity that belongs to every human being by the simple fact of existing" (§52) [1]. This is not a merely theological declaration; it is an affirmation that aims to give moral grounding to any technology that claims to validate the presence and integrity of an individual.

The question this chapter poses, as the thread running through the earlier sections on digital identity (ch. 47) and the crisis of truth brought on by deepfakes (ch. 10), is clear: if dignity is born simply of existing, what role can a tool that proves existence — t2p — play on this stage? Can this technology serve as a guardian of dignity, certifying that someone or something exists without reducing it to a mere set of data? Or could it, on the contrary, become the "technological power with a private face" the encyclical denounces (§5) [1], transforming the person into a commodity usable by economic and political interests?

This intersection between faith and science, between morality and innovation, sets the frame in which both the promise and the peril of any mechanism claiming to prove existence must be weighed. The answer will determine whether "Magnifica Humanitas" becomes an ethical compass capable of driving a future where technology not only recognizes but also protects the intrinsic dignity of every human being.

A believer, an atheist and the same dignity

That the Vatican chose Chris Olah to present the text beside the pope was no decorative gesture. Olah was not there as a discreet believer invited out of courtesy, but as one of the engineers who build, from the inside, the models the world is only beginning to understand. And his gaze — precisely because it comes from outside the faith — gives this chapter the tone we want to prevail in it: humanity first, religion second.

His words at the Vatican could be the motto of this book. He praised the Church as an "informed critic" and spoke of the beginning of "a long collaboration between those of us who are building this and those who can see what we, from inside, cannot" [7]. A single sentence held the entire thesis of *time2act*: whoever makes the technology cannot judge themselves. That is why Olah called for dialogue with people "not motivated by profit" [7] and warned, bluntly, that the industry "can't govern itself" [8]. It is exactly the

reason for being of an external, verifiable proof: that the truth not depend on the good faith of whoever produces it.

This convergence — a pope and an atheist affirming the same thing — is, in itself, the demonstration that the dignity of those who exist is not the doctrine of a few, but a minimum common humanity. When Leo XIV writes that every person is "unique and irreplaceable" [1], he does not ask us to believe it on faith: he names an evident truth that a skeptical engineer recognizes just as well. And it is on this universal minimum — not on any particular theology — that a tool like the proof of existence (t2p) can aspire to serve everyone. Because the same dignity belongs to us all, believers and non-believers, by the sole fact of existing; and a technology that limits itself to bearing witness to it, without claiming to grant it or measure it, can belong to everyone while belonging to no one.

One hundred and thirty-five years after *Rerum Novarum*

From the moment Pope Leo XIII published *Rerum Novarum* (1891), Catholic social doctrine has set the rhythm of a worldwide conversation on labor, dignity and collective responsibility [2]. That document, born of the tensions of the Industrial Revolution, laid down as its cornerstone the dignity of the worker in the face of exploitation. A century later, *Centesimus Annus* (1991) celebrated the centenary of this mission, reminding us that economic progress must be at the service of human well-being and not the other way around [5]. In 2015, with *Laudato Si'*, Francis widened this frame to the common home, warning of the danger of a technocratic paradigm that reduced nature — and, by extension, the person — to mere exploitable resources [3].

In 2026, in May, Pope Leo XIV published his first encyclical, *Magnifica Humanitas*, a work that positions itself as the direct continuation of this legacy. Signed on the 15th of the month and released at its end, to coincide with the anniversary of *Rerum Novarum* [1], it opens with the declaration that resounds at the heart of every debate: "the dignity that belongs to every human being by the simple fact of existing" (per the corresponding paragraph) [1]. This affirmation, simple but profound, establishes the moral base of the *time2act* project: a technology — the proof of existence (t2p) — that can validate the presence and integrity of a person without attempting to assign a dignity that is already intrinsic.

The key element of the encyclical is found in the first paragraph, where the pope poses a decisive choice for humanity: "a decisive choice: to raise a new tower of Babel or to build the city where God and humanity dwell together" (§1) [1]. This metaphor is not mere rhetoric; it represents the turning point at which emerging technologies, such as artificial intelligence and the blockchain, can become either bridges toward a more just coexistence or new towers that separate and fragment. The risk, already flagged in paragraph §9 of the encyclical, is that "technology can heal, connect, educate... but it can also divide, discard" [1]; and further still, that technological power take on a private, dominant face (§5) [1].

As we have seen in chapters 10, 47 and 112, this narrative converges on a moral choice: if the proof of existence is developed as a humble tool, capable only of certifying that someone exists and is whole, it can serve as a bastion for human dignity so long as it does not become the denounced "technological power with a private face" (§5) [1]. *Magnifica Humanitas* reminds us that the future depends on our ethical commitment; our task, as readers and creators of these technologies, is to choose the "city where God and humanity dwell together," making sure that every innovation is a bridge toward dignity and not a new tower of Babel that separates us.

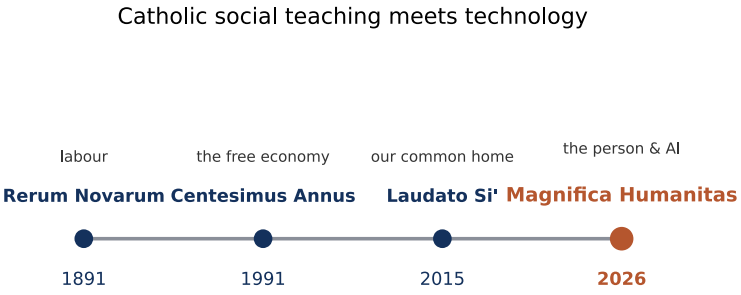


Figure 114.1 — The Church's social doctrine facing the economy and technology: from 'Rerum Novarum' (1891, labor) to 'Laudato Si" (2015, the common home) and 'Magnifica Humanitas' (2026, the person in the age of AI).

The dignity of those who exist

On May 25, 2026 Pope Leo XIV personally presented his first encyclical, *Magnifica Humanitas*, which becomes the moral lighthouse of this book. In paragraph fifty-two the encyclical affirms with clarity and force: "the dignity that belongs to every human being by the simple fact of existing" (§52) [1]. This affirmation, rooted in the Church's social tradition since *Rerum Novarum* (1891) — invoked by Leo XIII as the foundation of the worker's dignity — [2], is not a mere abstract declaration; it is the starting point for any technology that aims to be responsible. Dignity, according to the encyclical, is prior and intrinsic: the person already possesses it simply by being alive, as paragraph fifty recalls: "man and woman are created in the image and likeness of the triune God" (§50) [1].

This moral premise resonates directly with the problem laid out in chapter 47 on digital identity. Millions of individuals around the world still have no recognized legal identity; their existence is reduced to a file number or, worse, to the total absence of any official record. In this context, *t2p* — the proof of existence we have developed — does not claim to confer dignity; its role is to certify a reality that already exists but is often denied by legal and administrative structures.

The risk of confusing certification with the conferral of dignity is flagged in the encyclical itself. In paragraph fifty-one, Leo XIV denounces the reduction of the person to a mere means: "the person ends up reduced to a means for obtaining results, a resource to be used and exploited" (§51) [1]. This is the danger the deepfakes described in chapter 10 bring to mind; without a reliable proof, identity can be replaced or erased by algorithms that put deception before transparency.

t2p, in this frame, acts as a shield against this erosion. When a system denies the legal existence of an individual — for example, a person without civil registration — the proof of existence can record immutably that that voice is theirs and that that event took place. This offers more than a legal defense; it is a moral vindication that supports the pope's affirmation: dignity is already present, it only needs to be recognized.

Yet the encyclical also warns us about private technological power (paragraph five): "technological power thus acquires an unprecedented, predominantly 'private' face" (§5) [1]. If the proof of existence becomes a tool controlled by entities with exclusive interests, its emancipatory potential is nullified and it turns into a new form of domination. Our responsibility, as creators and users, is to guarantee that *t2p* be an open, transparent infrastructure subject to principles of subsidiarity (paragraph sixty-eight): "what people can do... must not be absorbed by higher instances" (§68) [1].

Ultimately, the proof of existence is a humble but powerful tool. It does not grant dignity — that already exists by the simple fact of being — but strengthens it by providing unequivocal testimony of our presence in the digital world. Used with wisdom and responsibility, it can be the bridge the encyclical describes between technology and the loving guardianship of humanity: a city where God and man dwell together, rather than a tower of Babel that separates us.

So that the person is not reduced to a data point

The encyclical clearly recalls the danger that threatens our essence: "the person ends up reduced to a means for obtaining results, a resource to be used and exploited" (as we saw in paragraph 51) [1]. This warning resonates especially in the era of digital identities, where the value of a life is often measured in bits, tokens or algorithmic scores. A badly designed system can turn us into mere records, fingerprints or profiles that technology treats as tradable resources, and not as beings with intrinsic dignity.

The proof-of-existence model (t2p) aims to be the direct opposite of this reduction. It is not a tool that assigns a value or a moral score; its goal is to certify, with minimal intervention, that someone exists, and to do so without creating a centralized file that can be exploited. The key to the design lies in three principles: **minimization of the fingerprint**, **locality of processing** and **personal ownership of the seal**.

1. **Minimize the fingerprint, not the content.** Instead of storing the whole history of interaction or complete biometric data, the protocol generates a unique hash that represents only verifiable presence. This translates into a small string of characters which, without the holder's private key,

cannot be recognized as personal information. This reduction prevents platforms from amassing large volumes of sensitive information and blocks later use for commercial or surveillance purposes.

2. **Processing on the device, not in central databases.** Verification takes place entirely on the user's terminal: their smartphone or an authorized IoT device. The hash computation and the digital signature are executed locally, without any external server having access to the original data. The information thus never leaves the environment controlled by the person, and the risk of a massive leak is drastically reduced.

3. **The seal belongs to the holder.** Once the proof is generated, it is stored in a digital wallet that only the owner can manage. The private key is unique and non-transferable; no external entity can revoke or modify it without the subject's explicit consent. This guarantees that the seal does not become a tradable asset, but a humble certification of presence.

These principles reflect the principle of subsidiarity indicated in the encyclical (§68) [1] — "what people can do... must not be absorbed by higher instances" — and, as we have seen before, they guarantee each individual total control over their own proof, protecting it from the temptation of private technological power described in the paragraph (§5) [1]. This is not only a technical matter, but an ethical commitment: technology must be a tool that **guards** and not a machine that **exploits**.

In short, the design of t2p answers Pope Leo XIV's warning, becoming a bridge between innovation and the loving guardianship of humanity. Applied with rigor and responsibility, it can prevent the person from being reduced to a data point, keeping alive the dignity that belongs to them by the simple fact of existing.

The private face of power

In May 2026 Pope Leo XIV personally presented his first encyclical, *Magnifica Humanitas*, with a clear moral mission: "the guardianship of the human person in the age of artificial intelligence" [4]. This moment, as we have seen, marked a symbolic anniversary and established the guiding thread of our book: the intrinsic dignity of every human being by the simple fact of existing, as the encyclical affirms (§52): "the dignity that belongs to every human being by the simple fact of existing" [1].

The "critique of the private face of power" unfolds across several paragraphs of *Magnifica Humanitas*. First, the document warns that "technological power thus acquires an unprecedented, predominantly 'private' face" and that we must ask "who holds this power today and toward what ends they direct it" (§5) [1]. This question becomes the cornerstone of democratic governance: if digital infrastructures — patents, algorithms, platforms, data — are controlled by a few actors, the risk is that technology stops being a tool of guardianship and turns into an extension of personalized power. Chapter 112, devoted to who controls the infrastructure, had already shown how the big platforms define conditions of access, rules of visibility, forms of relationship and even economic opportunities (§71) [1]; thus what seemed a technical analysis reveals itself as a direct continuation of the papal critique.

Next, the encyclical stresses that "technique takes on the face of whoever conceives it, finances it, regulates it, uses it" (§9) [1]. Every deep-learning model is an extension of its creator; when this responsibility falls into private hands, the person is reduced to a mere resource. The encyclical's warning — "the person ends up reduced to a means for obtaining results, a resource to be used and exploited" (§51) [1] — resonates with the conclusions of chapter 10 on the crisis of truth and deepfakes; as we have seen, the absence of a reliable proof facilitates the mass manipulation of images and voices.

In this context, *t2p* appears as a humble but potent technical answer: it certifies that someone or something exists and is whole without granting it dignity. Dignity is already guaranteed by the simple fact of existing (§52); the proof can only reinforce this reality against deception. If the mechanism is put at the service of a small group to consolidate power — like a "digital letter of credit" that controls access to essential services — it becomes the very private face the doctrine condemns.

Thus the intersection between the encyclical and chapter 112 shows that any implementation of *t2p* must be subordinated to principles of subsidiarity (§68) and to an open governance ensuring that "what people can do... must not be absorbed by higher instances" [1]. Only then can technology fulfill its mission of guardianship, as the pope recalls in §15: "to guard with love this magnificent humanity... which no machine will ever be able to replace." Our task is to translate this love into democratic, transparent structures, preventing innovation from becoming an instrument of domination.

The universal destination of digital goods

Pope Leo XIV recalls with surprising clarity that the "universal destination of goods today includes 'patents, algorithms, digital platforms, technological infrastructures, data'" (§67) [1]. This moral warning about the new geography of power remains in force: blockchain networks and artificial-intelligence systems shape the architecture of trust, and what used to be intellectual property becomes an integral part of the social fabric. Every token and every model trained on massive data become collective goods that affect the dignity of those who exist.

According to the encyclical, the ethical answer is the principle of subsidiarity: "what people can do... must not be absorbed by higher instances" (§68) [1]. In technical terms, this implies a proof-of-existence infrastructure (t2p) that is fully open and common. It cannot be a service managed by a corporation which, as paragraph §5 warns, has acquired "an unprecedented, predominantly 'private' face" of technological power (§5) [1]. Nominal centralization — where ownership of a node or a key is concentrated in few hands — does not guarantee the transparency or the accountability that human dignity demands.

In practice, this means a genuinely decentralized architecture. Distributed networks must allow any user to verify their own identity and that of others without depending on a higher authority. As we have seen with the platforms that define conditions of access (§71) [1], the t2p system should establish open rules, independent audits and mechanisms of participatory governance. A model of "governance by consortium," in which decisions on protocol upgrades are taken by vote among public entities and civil-society organizations, is one possible frame for avoiding capture by a single actor.

This configuration carries a double promise. First, it serves the dignity of the excluded: people living in contexts of vulnerability — digital refugees, informal-sector workers or victims of deepfakes — can prove their existence and protect their rights without depending on a company that might sell their data. Second, it reinforces truth against deception: a verifiable proof of existence prevents anyone from manipulating images or audio to discredit an individual, as we have seen in chapter 10.

Nevertheless, the danger persists. If the infrastructure falls under the control of a private group, it becomes "a resource to be used and exploited" (§51) [1], reversing the goal of guarding with love this magnificent humanity (§15) [1]. The challenge, therefore, is not only technological but institutional: legal frameworks must be established that mandate transparency and accountability, just as the social doctrine of 'Rerum Novarum' (1891) demanded in its day for the industrial worker [2].

In short, the universal destination of digital goods must align with subsidiarity and genuine decentralization. Only then can the t2p tool be a "humble tool" that certifies existence without usurping the intrinsic dignity of those who exist, becoming the true bridge between technology and morality that the encyclical so insisted on defending.

The limit no machine can replace

There is one last limit this book must acknowledge. The encyclical issues a call to responsibility: "the dignity that belongs to every human being by the simple fact of existing" (§52) and, at the same time, reminds us that this dignity cannot be delegated to any algorithm or any blockchain. This double affirmation sets the moral frame in which our technological proposal — t2p — is inscribed.

Humility is the central principle the encyclical asks of the creators and users of technology. Thus, "to accept the limits and fragility of humanity without considering them an error to be corrected" (§12) means recognizing that no system can capture the totality of the human being. The proof of existence, as described throughout the book, is a tool that certifies that a file existed intact and that a person is who they say they are. It does not claim to measure the soul, merit or virtue; it does not claim to replace human trust or the care that only a direct relationship can offer. This coincides with Leo XIV's warning about the danger of reducing the person to a means: "the person ends up reduced to a means for obtaining results, a resource to be used and exploited" (§51) [1].

t2p presents itself as a "humble tool." Its function is strictly documentary: it records the integrity of a file at the moment of its creation and verifies that the issuer is who they declare to be. This certification can protect the rights of the most vulnerable — like the victims of deepfakes described in chapter 10 — but it has no capacity to assign dignity or determine the intrinsic value of a life, as the pope affirms: "to guard with love this magnificent humanity... which no machine will ever be able to replace" (§15) [1].

The tool's limit is clear in its technical specifications. In the context of digital identity (chapter 47), t2p can guarantee that a digitized credential is authentic, but it does not replace the personal recognition of a face-to-face exchange. Similarly, the chapter on who controls the infrastructure (chapter 112) intimates that the centralization of technological power — "technological power thus acquires an unprecedented, predominantly 'private' face" (§5) — can turn any tool into an instrument of domination. Our proposal upholds the subsidiarity proclaimed by the encyclical: "what people can do... must not be absorbed by higher instances" (§68) [1].

Thus the final message of this section is a call to moderation. Technology can and must serve as a support for human dignity, but never as a substitute for the moral commitment that man must hold toward himself and toward others. In this balance — between the transformative potential of the proof of existence and respect for the intrinsic limitations of any machine — lies the true value of *t2p*: a tool at the service of the person, not an oracle of their worth.

The promise and the peril

In the ninth paragraph of *Magnifica Humanitas* Pope Leo XIV describes technology as double-faced: "technology can heal, connect, educate, care for the common home; but it can also divide, discard" (§9) [1]. This paradox is the guiding thread of what we call *t2p*: a tool capable of certifying the existence of a person or a digital asset with a degree of integrity until now unattainable. If the proof of existence is deployed as a mechanism of transparency, it can offer a concrete answer to the crises documented earlier — the manipulation of truth and the erosion of digital identity — and, as we have seen, turn the testimony of an existence into a public record that no one can deny or falsify. This is, literally, making truth verifiable again and giving "proven existence" to those who have been privately denied it.

Yet the same technological power that gives life to *t2p* contains an intrinsic danger. The encyclical warns that "technological power thus acquires an unprecedented, predominantly 'private' face" (§5) [1]; and that the person can end up reduced to a means for obtaining results, a resource to be used and exploited. If one entity — public or private — controls the issuance and recognition of these certificates, a new form of domination is created: deciding who is "existent" and who is not. In that scenario, the proof of existence becomes an instrument of mass surveillance, similar to the one we have observed, and could be used to establish a "censorship of existence," excluding individuals or groups from economic and social life simply because they lack the right certificate.

The encyclical offers a powerful metaphor: "a decisive choice: to raise a new tower of Babel or to build the city where God and humanity dwell together" (§1) [1]. If the proof of existence is governed by private interests, a digital Tower of Babel is built — a labyrinth of incompatible protocols that fragments human cooperation. If, instead, its development follows the principles of subsidiarity, as we have seen, and of loving guardianship (§15) [1], a "holy city" can arise where every certificate is testimony to a pre-existing dignity, not its source. The danger, then, lies not in the tool itself, but in the intent and the structures of government that surround it.

The reader must take up this discernment: recognizing that *t2p* can be a key to restoring trust in a world saturated with deepfakes and false identities, but that it is also liable to reinforce hierarchies of power if left unsupervised. Our task — and the book's — is to offer not only the technology, but the ethical framework that guarantees that the proof of existence serves "to guard with love this magnificent humanity" (§15) [1], and not to turn it into a new form of domination. In this balance between promise and peril lies the true value of *t2p*: a humble instrument that certifies existence without attempting to grant dignity, leaving the latter as the intrinsic prerogative of every human being.

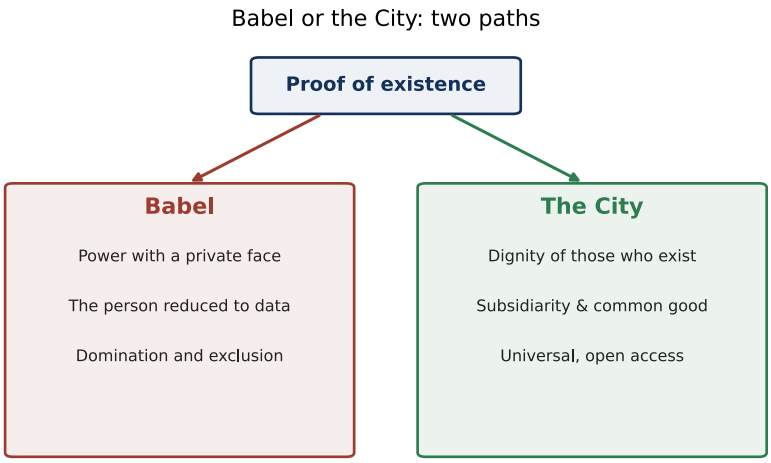


Figure 114.2 — The same proof of existence, two paths (§1): toward "Babel" (private power, the person reduced to data, domination) or toward the "city" (the dignity of those who exist, subsidiarity, universal access). The difference is not technical, but one of governance and purpose.

It's time to act

The warehouse of decisions that artificial intelligence has already made about our daily lives stretches out like an endless carpet, and at its center there is a clear call: "Let us choose... the path of Nehemiah... transforming diversity into a resource" (§10) [1]. These words are not just a moral ideal; they are a plan of action which, through the principles matured throughout *time2act*, can guide us toward an era where technology truly serves human dignity.

First principle — Proof of existence for the excluded. As we have seen in chapter 47, digital identity becomes a refuge for those who lack one, be it an undocumented migrant or a marginalized community living outside the traditional networks. The proof of existence (t2p) must be the bridge connecting these people with their legal and social recognition, without any entity being able to deny them access to fundamental rights.

Second principle — Truth as defense against deepfakes. The crisis of truth described in chapter 10 shows that manipulated content threatens to nullify any basis for authentic dialogue. In this context, the proof of existence acts as an immutable fingerprint that certifies the origin and integrity of statements, preventing information from becoming a mere visual illusion.

Third principle — The person must not be reduced to a data point. The encyclical recalls that "the person ends up reduced to a means for obtaining results, a resource to be used and exploited" (§51) [1]; as we have seen, the proof of existence must remain in the hands of its holder, not in centralized databases that could turn it into a tradable object. Every individual must be able to validate their own status without depending on a third party that controls them.

Fourth principle — Democratic governance against the private face. The danger of "technological power... with an unprecedented, predominantly 'private' face" (§5) [1] was laid out in chapter 112, which details how the concentration of infrastructure can erect a new digital aristocracy. The answer must be a model of open governance, with representative councils that reflect human diversity.

Fifth principle — Common infrastructure and subsidiarity. The encyclical stresses that "the universal destination of goods today includes patents, algorithms, digital platforms, technological infrastructures, data" (§67) [1] and that the principle of subsidiarity asks that "what people can do... must not be absorbed by higher instances" (§68) [1]; thus the infrastructure sustaining the proof of existence must be a public, cooperative, decentralized network, accessible to all without economic or political barriers.

In short, the path of Nehemiah invites us to transform diversity into a resource that guarantees identity for the invisible, truth for fragmented societies and autonomy for every person. This is the last step of the journey begun with *Rerum Novarum* and continued with *Laudato Si'*; now, at the heart of artificial intelligence, we face the most crucial decision: **to guard the magnificent humanity of those who exist — that, in the age of AI, is time2act**.

[1] Leo XIV — *Encyclical letter "Magnifica Humanitas"* (Vatican, May 15-25, 2026; on the guardianship of the human person in the age of artificial intelligence). https://www.vatican.va/content/leo-xiv/es/encyclicals/documents/20260515-magnifica-humanitas.html ↩

[2] Leo XIII — *Rerum Novarum* (1891; the origin of the Church's social doctrine). https://www.vatican.va/content/leo-xiii/es/encyclicals/documents/hf_l-xiii_enc_15051891_rerum-novarum.html ↩

[3] Francis — *Laudato Si'* (2015; the common home and the technocratic paradigm). https://www.vatican.va/content/francesco/es/encyclicals/documents/papa-francesco_20150524_enciclica-laudato-si.html ↩

[4] Vatican News — *Leo XIV presents his first encyclical "Magnifica Humanitas"* (May 25, 2026). https://www.vaticannews.va/es/papa/news/2026-05/papa-leon-xiv-primera-enciclica-magnifica-humanitas.html ↩

[5] John Paul II — *Centesimus Annus* (1991; centenary of 'Rerum Novarum'). https://www.vatican.va/content/john-paul-ii/es/encyclicals/documents/hf_jp-ii_enc_01051991_centesimus-annus.html ↩

[6] Vatican News / The Washington Post — *Leo XIV presents "Magnifica Humanitas" in person, with Anthropic cofounder Chris Olah; "with our differences, we can listen to one another"* (25-05-2026). https://www.vaticannews.va/en/pope/news/2026-05/pope-leo-xiv-encyclical-magnifica-humanitas-ai.html ↩

[7] Anthropic — *Chris Olah's remarks on Pope Leo XIV's encyclical "Magnifica Humanitas"* (the Church as "informed critics"; "a long collaboration between those of us who are building this and those who can see what we, from inside, cannot"; dialogue with people "not motivated by profit") (2026). https://www.anthropic.com/news/chris-olah-pope-leo-encyclical ↩

[8] Fortune — *Who is Chris Olah? The atheist Anthropic cofounder the Pope chose to sit beside him at the Vatican* (raised evangelical in Toronto, an atheist since 15; the industry "can't govern itself") (03-06-2026). https://fortune.com/2026/06/03/who-chris-olah-anthropic-cofounder-atheist-pope-leo-vatican-ai/ ↩

Bibliography

1. IEA, *Oil 2025 — Executive summary*. https://www.iea.org/reports/oil-2025/executive-summary
2. IRENA (via pv magazine), *Global average solar LCOE $0.043/kWh in 2024*. https://www.pv-magazine.com/2025/07/23/global-average-solar-lcoe-stood-at-0-043-kwh-in-2024-says-irena/
3. Copernicus C3S, *2024, first year above 1.5 °C*. https://climate.copernicus.eu/copernicus-2024-first-year-exceed-15degc-above-pre-industrial-level
4. IEA, *Energy and AI — Executive summary*. https://www.iea.org/reports/energy-and-ai/executive-summary
5. ethereum.org, *The Merge*. https://ethereum.org/roadmap/merge/
6. C2PA (2026), *how the provenance standard works and its limits*. https://truescreen.io/articles/c2pa-standard-history-limitations/
7. Tapscott, Don and Tapscott, Alex — *Blockchain Revolution* (Portfolio/Penguin, 2016). https://blockchain-revolution.com/
8. Mougayar, William — *Explaining the Blockchain via a Google Docs Analogy* (2016). https://medium.com/@wmougayar/explaining-the-blockchain-via-a-google-docs-analogy-56326b1bbe58
9. EU AI Act — *Article 50: Transparency Obligations* (Regulation (EU) 2024/1689). https://artificialintelligenceact.eu/article/50/
10. European Commission — *European Digital Identity Wallet (eIDAS 2.0)*. https://www.eudi-wallet.eu/
11. C2PA — *Content Credentials*. https://c2pa.org/
12. Deloitte — *Generative AI and the fraud risk for financial services* (AI fraud projection). https://www2.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html
13. Europol — *Facing reality? Law enforcement and the challenge of deepfakes*. https://www.europol.europa.eu/publications-events/publications/facing-reality-law-enforcement-and-challenge-of-deepfakes
14. t2p — *User Manual (Data Integrity Solution)*, v1.0, 2026.
15. IEA, *Global Energy Review 2025 — Key findings*. https://www.iea.org/reports/global-energy-review-2025/key-findings
16. EIA/IEA, *Brent prices and OPEC+ (2022-2025)*. https://www.eia.gov/todayinenergy/detail.php?id=55020
17. IEA, *Oil Market Report (Sept. 2025) — US shale*. https://www.iea.org/reports/oil-market-report-september-2025
18. BloombergNEF / IEA, *Global EV Outlook 2025 and battery prices*. https://www.iea.org/reports/global-ev-outlook-2025/electric-vehicle-batteries
19. IEA / CFR, *Chinese dominance in minerals and clean technologies (2024-2025)*. https://www.iea.org/commentaries/with-new-export-controls-on-critical-minerals-supply-concentration-risks-become-reality
20. Sightline / analysis, *Chinese clean exports to the Global South (2024)*. https://www.sightlineclimate.com/research/a-tour-of-chinas-electrostate
21. IPCC AR6 (via COP28 Global Stocktake), *emissions cuts 2030/2035*. https://www.iea.org/topics/cop28-tracking-the-energy-outcomes
22. UNFCCC / European Commission, *outcomes of COP28-COP30*. https://climate.ec.europa.eu/news-other-reads/news/what-did-cop30-achieve-2025-12-01_en
23. IEA, *Energy Efficiency 2024 — Executive summary*. https://www.iea.org/reports/energy-efficiency-2024/executive-summary
24. *Jevons paradox / rebound effect* (review). https://en.wikipedia.org/wiki/Jevons_paradox
25. Reclaim Finance, *Assessment of oil & gas companies' climate strategy (2024)*. https://reclaimfinance.org/site/en/assessment-of-oil-and-gas-companies-climate-strategy/
26. The Bureau of Investigative Journalism, *BP rowing back on renewables (2025)*. https://www.thebureauinvestigates.com/stories/2025-03-06/bp-renewables-net-zero-banks
27. IISD, *Public money: 5x more to fossil fuels than renewables*. https://www.iisd.org/articles/press-release/governments-five-times-more-public-money-fossil-fuels-than-renewables
28. Google DeepMind, *AI reduces data centre cooling energy by 40%*. https://deepmind.google/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/

29. Union of Concerned Scientists, *Fossil fuel industry's decades of deceit*. https://www.ucs.org/about/news/new-ucs-report-details-fossil-fuel-industry-decades-deceit
30. United Nations, *Greenwashing*. https://www.un.org/en/climatechange/science/climate-issues/greenwashing
31. (deepfakes) Blockchain Council / Europol — see research 00 §6. https://www.blockchain-council.org/blockchain/preventing-deepfake-fraud-with-blockchain-content-provenance-verification-workflows/
32. Beth Simone Noveck — *Reboot: AI and the Race to Save Democracy* (Yale UP, 2026), ch. 3.
33. FCC (2024). https://www.fcc.gov/document/fcc-makes-ai-generated-voice-robocalls-illegal
34. Google DeepMind / Data Center Frontier, *AI for the grid*. https://www.datacenterfrontier.com/energy/article/55282719/googles-ai-powered-grid-revolution-how-data-centers-are-reshaping-the-us-power-landscape
35. IAEA / NucNet, *advanced nuclear and SMRs for AI demand (2024-2025)*. https://www.iaea.org/bulletin/data-centres-artificial-intelligence-and-cryptocurrencies-eye-advanced-nuclear-to-meet-growing-power-needs
36. RWA.io / Powerledger, *tokenized RECs and MRV on blockchain*. https://www.rwa.io/post/renewable-energy-credits-on-blockchain-a-tokenized-approach
37. Frontiers in Blockchain (2025), *state of blockchain carbon markets*. https://www.frontiersin.org/journals/blockchain/articles/10.3389/fbloc.2025.1603695/full
38. OpenTimestamps — *a timestamping proof standard*. https://opentimestamps.org/
39. London Blockchain — *Can blockchain save truth? (2025)*. https://londonblockchain.net/blog/blockchain-in-action/can-blockchain-save-truth-tackling-deepfakes-and-disinformation-in-2025/
40. C2PA — *Content Credentials Explainer (official specification)*. https://spec.c2pa.org/specifications/specifications/2.4/explainer/Explainer.html
41. arXiv — *Provenance Verification via a Perceptual Hash Registry Anchored on Blockchain*. https://arxiv.org/html/2602.02412v1
42. Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 3.
43. IEA, *Global Hydrogen Review 2025* / CleanTechnica. https://www.iea.org/reports/global-hydrogen-review-2025/executive-summary
44. IEA / The Extractor, *green H_2 cost 2022-2025*. https://cleantechnica.com/2025/07/07/hydrogens-brutal-month-billions-lost-as-mega-projects-collapse/
45. arXiv — *DeepFake-Eval-2024: A Multi-Modal In-the-Wild Benchmark of Deepfakes*. https://arxiv.org/html/2503.02857v4
46. ManageEngine Insights — *Content provenance is our best chance in the fight against deepfakes* (Hany Farid). https://insights.manageengine.com/artificial-intelligence/content-provenance-is-our-best-chance-in-the-fight-against-deepfakes/
47. WITNESS — *Prepare, Don't Panic: Synthetic Media and Deepfakes* (Sam Gregory). https://www.witness.org/portfolio_page/sam-gregory/
48. CISA / NSA / FBI — *Contextualizing Deepfake Threats to Organizations* (2023). https://www.cisa.gov/news-events/alerts/2023/09/12/nsa-fbi-and-cisa-release-cybersecurity-information-sheet-deepfake-threats
49. C2PA — *Content Credentials Explainer*. https://spec.c2pa.org/specifications/specifications/2.2/explainer/_attachments/Explainer.pdf
50. Peter Todd — *Carbon Dating the Internet Archive with OpenTimestamps*. https://petertodd.org/2017/carbon-dating-the-internet-archive-with-opentimestamps
51. ProofStamper — *File timestamp tools compared (RFC 3161 + blockchain anchoring)*. https://proofstamper.com/en/compare/file-timestamp-tools-2026/
52. Sumsub — *Global deepfake incidents surge tenfold*. https://sumsub.com/newsroom/sumsub-research-global-deepfake-incidents-surge-tenfold-from-2022-to-2023/
53. CNN Business — *Finance worker pays out $25 million after video call with deepfake 'CFO'* (Arup case). https://www.cnn.com/2024/05/16/tech/arup-deepfake-scam-loss-hong-kong-intl-hnk · World Economic Forum. https://www.weforum.org/stories/2025/02/deepfake-ai-cybercrime-arup/
54. Deloitte — *Generative AI is expected to magnify the risk of deepfakes and other fraud in banking* (≈$40 billion in 2027). https://www.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html
55. FinCEN — *Alert on Fraud Schemes Involving Deepfake Media* (Nov. 13, 2024). https://www.fincen.gov/system/files/shared/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf

56. iProov / Biometric Update — *Biometric injection attack surge* (KYC, NIST SP 800-63-4). https://www.biometricupdate.com/202604/biometric-injection-attack-surge-spreads-to-ios-iproov-report
57. Signaturit — *eIDAS 2.0: Qualified Electronic Ledgers and Archiving*. https://www.signaturit.com/blog/eidas-2-regulation/
58. Thomson Reuters Institute — *Deepfakes on trial: evidence authentication (FRE 707, 901(c))*. https://www.thomsonreuters.com/en-us/posts/ai-in-courts/deepfakes-evidence-authentication/
59. Forbes (Thomas Brewster) — *Fraudsters Cloned Company Director's Voice In $35 Million Heist, Police Find* (14-10-2021; UAE court documents: at least 17 people involved, $400,000 traced to Centennial Bank accounts in the US). https://www.forbes.com/sites/thomasbrewster/2021/10/14/huge-bank-fraud-uses-deep-fake-voice-tech-to-steal-millions/
60. NPR — *Fake viral images of an explosion at the Pentagon were probably created by AI* (May 2023). https://www.npr.org/2023/05/22/1177590231/fake-viral-images-of-an-explosion-at-the-pentagon-were-probably-created-by-ai
61. NPR / Harvard Ash Center — *AI was everywhere in 2024's elections, but deepfakes were only part of the picture*. https://ash.harvard.edu/articles/the-apocalypse-that-wasnt-ai-was-everywhere-in-2024s-elections-but-deepfakes-and-misinformation-were-only-part-of-the-picture/
62. Content Authenticity Initiative — *5,000 members* (2025). https://contentauthenticity.org/blog/5000-members-building-momentum-for-a-more-trustworthy-digital-world
63. Nikon Rumors — *AFP validates a C2PA certification test using a prototype Nikon camera* (Feb 2025). https://nikonrumors.com/2025/02/18/afp-validates-a-certification-test-of-its-photos-via-the-c2pa-standard-using-a-prototype-nikon-camera.aspx/
64. IPTC — *C2PA news: France Télévisions, IBC "Stamping Your Content"*. https://iptc.org/news/tag/c2pa/
65. Content Authenticity Initiative — *Leica launches world's first camera with Content Credentials*. https://contentauthenticity.org/blog/leica-launches-worlds-first-camera-with-content-credentials
66. arXiv — *Why the C2PA specifications fall short* (strippable metadata). https://arxiv.org/html/2604.24890v1
67. NSA/CISA — *Content Credentials: Strengthening Multimedia Integrity* (Jan 2025); JPEG Trust ISO/IEC 21617-1. https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF
68. Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 3; Infobae (2025). https://www.infobae.com/peru/2025/06/06/blockchain-y-elecciones-2026-como-se-desarrollara-en-peru-el-voto-digital-obligatorio-para-militares-y-ciudadanos-en-el-extranjero/
69. Mirsky et al. — *CT-GAN: Malicious Tampering of 3D Medical Imagery* (USENIX Security 2019). https://www.usenix.org/conference/usenixsecurity19/presentation/mirsky · summary with figures: https://www.pindrop.com/article/threat-deepfakes-in-healthcare/
70. TruthScan / Pindrop — *AI-driven fraud in healthcare: deepfake doctors and synthetic patients* (2025). https://truthscan.com/blog/ai-driven-fraud-in-global-healthcare-2025-trends-and-countermeasures/
71. STAT News — *Cassava Sciences scientist indicted (fabricated Western Blots)* (28-Jun-2024); Retraction Watch / Bik. https://www.statnews.com/2024/06/28/cassava-sciences-indictment-alzheimers-hoau-yan-wang/
72. FDA — *Cybersecurity in Medical Devices: final guidance* (Federal Register, 27-Jun-2025). https://www.govinfo.gov/content/pkg/FR-2025-06-27/html/2025-11669.htm
73. HIPAA Vault — *Blockchain integration for healthcare records (audit trails)* (2025). https://www.hipaavault.com/resources/blockchain-integration-healthcare-records/
74. Bentham/RRCT (PubMed) — *Blockchain as a Prime Guardian of clinical-trial data* (2025); Triall. https://pubmed.ncbi.nlm.nih.gov/39851126/
75. HEPI — *Student Generative AI Survey 2025*. https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/
76. TEQSA — *The evolving risk to academic integrity posed by generative AI* (2024). https://www.teqsa.gov.au/sites/default/files/2024-08/evolving-risk-to-academic-integrity-posed-by-generative-artificial-intelligence.pdf
77. U.S. DOJ — *Operation Nightingale: fraudulent nursing diplomas* (2023; Phase II 2025). https://www.justice.gov/usao-sdfl/pr/fraudulent-nursing-diploma-scheme-leads-federal-convictions

78. Open edX — *Proctoring in the age of AI: deepfake impersonation* (2025). https://openedx.org/blog/proctoring-in-the-age-of-ai-tackling-modern-cheating-techniques/
79. MIT Technology Review — *Digital diplomas (Blockcerts)* (2018); MIT Registrar. https://www.technologyreview.com/2018/04/25/143311/digital-diplomas/
80. European Commission — *EBSI Verifiable Credentials*; Europass *European Digital Credentials for Learning*; W3C VC Data Model 2.0. https://europass.europa.eu/en/european-digital-credentials-learning
81. NPR — *Zelenskyy deepfake telling Ukrainians to surrender* (16-Mar-2022). https://www.npr.org/2022/03/16/1087062648/deepfake-video-zelenskyy-experts-war-manipulation-ukraine-russia
82. EU DisinfoLab — *Doppelganger* (Russian disinformation campaign). https://www.disinfo.eu/doppelganger-hub/
83. FCC — *Proposed $6M fine for AI-deepfake Biden robocall* (2024); IPI (Slovakia). https://www.fcc.gov/document/fcc-proposes-6m-fine-nh-robocalls
84. University of Washington — *Deepfake geography: AI falsifies satellite images* (Bo Zhao, 2021). https://www.washington.edu/news/2021/04/21/a-growing-problem-of-deepfake-geography-how-ai-falsifies-satellite-images/
85. International Bar Association — *eyeWitness to Atrocities* (hash at capture; cases in Ukraine). https://www.ibanet.org/IBA-founded-eyeWitness-to-Atrocities-marks-ten-years-of-capturing-evidence-of-human-rights-crimes
86. USC Shoah Foundation — *Starling Lab & Hala Systems file cryptographic war-crimes evidence to the ICC* (2022); Berkeley Protocol (OHCHR). https://sfi.usc.edu/news/2022/06/33571-starling-lab-and-hala-systems-file-cryptographic-submission-evidence-war-crimes
87. DARPA — *Semantic Forensics (SemaFor)* and *Deepfake Defense* (2024); NATO StratCom COE. https://www.darpa.mil/research/programs/semantic-forensics
88. European Commission — *About the EU Emissions Trading System* (coverage ~40%, €175 billion raised, −47% vs 2005, −62% by 2030). https://climate.ec.europa.eu/eu-action/carbon-markets/about-eu-ets_en
89. Trading Economics — *EU Carbon Permits* (EUA €80.58/t on 19-Jun-2026; all-time high €105.73 in Feb-2023). https://tradingeconomics.com/commodity/carbon
90. European Commission (Taxation and Customs Union) — *Carbon Border Adjustment Mechanism (CBAM)* (definitive phase from 1-Jan-2026). https://taxation-customs.ec.europa.eu/carbon-border-adjustment-mechanism_en
91. S&P Global — *Brussels confirms first CBAM certificate price for Q1 2026 at €75.36/tCO2e* (7-Apr-2026). https://www.spglobal.com/energy/en/news-research/latest-news/energy-transition/040726-brussels-confirms-first-cbam-certificate-price-for-q1-2026-at-eur7536mtco2e
92. International Carbon Action Partnership (ICAP) — *EU Emissions Trading System (EU ETS)* (MSR, CBAM factor, free allocation). https://icapcarbonaction.com/en/ets/eu-emissions-trading-system-eu-ets
93. European Commission — *ETS2: buildings, road transport and additional sectors* (start postponed to 2028; agreement Nov-2025). https://climate.ec.europa.eu/eu-action/carbon-markets/ets2-buildings-road-transport-and-additional-sectors_en
94. EASA — *ICAO CORSIA* (2024 international emissions ≈ 597 $MtCO_2$; ≈ 58 $MtCO_2$ to offset). https://www.easa.europa.eu/en/domains/environment/icao-carbon-offsetting-and-reduction-scheme-international-aviation-corsia
95. European Commission — *Monitoring, reporting and verification of EU ETS emissions* (MRR and AVR; verified report before March 31). https://climate.ec.europa.eu/eu-action/carbon-markets/eu-emissions-trading-system-eu-ets/monitoring-reporting-and-verification_en
96. Europol — *Carbon Credit fraud causes more than 5 billion euros damage for European Taxpayer* (up to 90% of volume fraudulent in some countries). https://www.europol.europa.eu/media-press/newsroom/news/carbon-credit-fraud-causes-more-5-billion-euros-damage-for-european-taxpayer
97. Marius-Christian Frunza — *Fraud and Carbon Markets: The Carbon Connection* (Routledge, 2013) and *An Assessment of Carousel VAT Fraud in the European Carbon Market* (Review of Law & Economics, 2014). https://www.routledge.com/Fraud-and-Carbon-Markets-The-Carbon-Connection/Frunza/p/book/9781138928091
98. European Commission — *Announcement: transitional measure for the EU ETS registry system* (19-Jan-2011) and OCCRP — *Germany: Six Executives Stand Trial for Emissions Fraud* (~€850 million; six convicted). https://climate.ec.europa.eu/news-your-voice/news/announcement-

transitional-measure-eu-ets-registry-system-2011-01-19_en · https://www.occrp.org/en/news/germany-six-executives-stand-trial-for-emissions-fraud
99. World Bank — *Why data infrastructure is key for a transparent carbon market* and the *Climate Action Data Trust* platform. https://blogs.worldbank.org/en/climatechange/why-data-infrastructure-key-transparent-carbon-market · https://climateactiondata.org/
100. The Guardian / Die Zeit / SourceMaterial — Patrick Greenfield, *Revealed: more than 90% of rainforest carbon offsets by biggest certifier are worthless* (18-Jan-2023). https://www.theguardian.com/environment/2023/jan/18/revealed-forest-carbon-offsets-biggest-provider-worthless-verra-aoe
101. Verra — *Verra Response to Guardian Article on Rainforest Offset Credits* (rebuttal). https://verra.org/verra-response-guardian-rainforest-carbon-offsets/
102. Mongabay — *Head of Verra, top carbon credit certifier, to leave in June* (resignation of David Antonioli, 22-May-2023). https://news.mongabay.com/2023/05/head-of-verra-top-carbon-credit-certifier-to-leave-in-june/
103. Verra — *Verra Acts on Kariba Project: Cancels Excess Credits and Advances Independent Review* (15,220,520 credits in excess out of 26,822,953; ~57%). https://verra.org/verra-acts-on-kariba-project-cancels-excess-credits-advances-independent-review/
104. South Pole — *Statement on the Kariba REDD+ project* (27-Oct-2023). https://www.southpole.com/news/statement-27october
105. Alejandro Guizar-Coutiño et al. — *A global evaluation of the effectiveness of voluntary REDD+ projects* (Conservation Biology, 2022; deforestation −47%). https://pmc.ncbi.nlm.nih.gov/articles/PMC10086997/
106. West et al. — *Action needed to make carbon offsets from forest conservation work for climate change mitigation* (Science, 2023; 26 sites in 6 countries; disputed result, arXiv reply 2312.06793). https://www.science.org/doi/10.1126/science.ade3535
107. Ecosystem Marketplace / Forest Trends — *State of the Voluntary Carbon Market 2024* (market ~$2 billion in 2021 → $723 million in 2023; volume −56%; average price $6.53/t). https://www.ecosystemmarketplace.com/articles/report-the-voluntary-carbon-market-contracted-in-2023-driven-by-drop-off-in-transactions-for-redd-and-renewable-energy/
108. ICVCM — *Core Carbon Principles* and *First high-integrity CCP-labelled carbon credits* (June 2024; ~27 million credits). https://icvcm.org/core-carbon-principles/ · https://icvcm.org/integrity-council-announces-first-high-integrity-ccp-labelled-carbon-credits-as-assessments-continue/
109. VCMI — *Claims Code of Practice* (28-June-2023). https://vcmintegrity.org/launch-claims-code/
110. CarbonPlan — Danny Cullenward and Grayson Badgley, *Zombies on the blockchain* (7-Apr-2022; ~28% zombie credits). https://carbonplan.org/research/toucan-crypto-offsets
111. Verra — *Verra Addresses Crypto Instruments and Tokens* (ban on tokenizing retired credits, 25-May-2022). https://verra.org/verra-addresses-crypto-instruments-and-tokens/
112. S&P Global Commodity Insights — *COP29: rules for UN-led carbon market under Article 6.4 approved in Baku* (11-Nov-2024). https://www.spglobal.com/commodity-insights/en/news-research/latest-news/energy-transition/111124-cop29-rules-for-un-led-carbon-market-under-article-64-approved-in-baku
113. CIFOR-ICRAF / Open Forest Protocol — *Blockchain platform offers reforestation efforts access to carbon finance* (open digital MRV). https://forestsnews.cifor.org/76905/blockchain-platform-offers-reforestation-efforts-access-to-carbon-finance
114. European Commission — *Energy Efficiency Directive* (Directive 2023/1791, Art. 8; new savings up to 1.9% annually in 2028-2030; "more than doubles" toward 2028). https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive_en
115. BOE — *Royal Decree 36/2023, of January 24* (CAE system; 1 CAE = 1 kWh of accredited final energy; obligated/delegated parties, verifier, coordinator). https://www.boe.es/buscar/doc.php?id=BOE-A-2023-2027
116. BOE — *Order TED/845/2023, of July 18* (catalog of standardized energy efficiency measures). https://www.boe.es/buscar/doc.php?id=BOE-A-2023-16940
117. Funcas — *Sistema de Certificados de Ahorro Energético (CAE): el motor que ya impulsa la transformación energética en España* (2024 results: ~1,200 measures, ~2 TWh, ~500,000 t CO_2, ~€230 million; industry 68%). https://www.funcas.es/articulos/sistema-de-certificados-de-ahorro-energetico-cae-el-motor-que-ya-impulsa-la-transformacion-energetica-en-espana/
118. MITECO — *Memoria de la Orden de Obligaciones del Sistema CAE 2025* (cumulative savings target 2021-2030: 53,593 ktoe). https://www.miteco.gob.es/content/dam/miteco/es/energia/files-1/es-

ES/Participacion/Documents/aeip-fnee-2025/Memoria%20OM_Obligaciones%202025_TA.pdf
119. Légifrance — *Décret n° 2025-1048 du 30 octobre 2025* (CEE, 6th period 2026-2030; 1,050 TWh cumac/year). https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000052486193
120. Légifrance — *Loi n° 2025-594 du 30 juin 2025* (penalties up to 10-12%; photographs *horodatées et géolocalisées* or remote video, kept 6 years). https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000051824277
121. DGCCRF (economie.gouv.fr) — *Rénovation énergétique : focus sur l'isolation à 1 euro* (469 establishments inspected; 56% non-compliance). https://www.economie.gouv.fr/dgccrf/laction-de-la-dgccrf/les-enquetes/renovation-energetique-focus-sur-lisolation-1-euro
122. Hellio / French government — *Fraude aux CEE* (€480 million of fraud blocked in 2024, €236 million in CEE; control rate 10%→15%). https://www.hellio.com/actualites/reglementation/fraude-cee
123. ANSA / PordenoneToday — *Contributi per 10 milioni su certificati bianchi inesistenti* (292 fictitious projects; 50,322 TEE; damage ~€10.4 million). https://www.ansa.it/sito/notizie/cronaca/2025/06/23/contributi-per-10-milioni-su-certificati-bianchi-inesistenti_dd57abfb-77ba-483a-8f42-c240874fc308.html
124. EnergiaMercato / Consiglio di Stato — *GSE and the annulment of energy efficiency titles* (the "progressive formation" doctrine; 4ESCO case). https://www.energiamercato.it/notizie/project-finance/gse-conferma-la-legittimita-dellannullamento-di-titoli-di-efficienza-energetica-a-4esco
125. Louis-Gaëtan Giraudet, Matthieu Glachant and Jean-Philippe Nicolaï — *Selling and Saving Energy: Energy Efficiency Obligations in Liberalized Energy Markets* (The Energy Journal, 2020). https://ideas.repec.org/f/pgi337.html
126. UNEF — *El autoconsumo fotovoltaico en España alcanza los 9,3 GW instalados* (report on 2025: +1,139 MW in the year, −3.7% vs 2024; 36,330 new households +229 MW; commercial +176 MW; industrial +679 MW; fifth generation source by installed capacity, ahead of nuclear; PNIEC target 19 GW); UNEF — *El autoconsumo sigue avanzando, superando ya el umbral de los 8 GW* (2024: +1,182 MW, cumulative 8,137 MW). https://www.unef.es/es/comunicacion/comunicacion-post/el-autoconsumo-fotovoltaico-en-espana-alcanza-los-9-3-gw-instalados · https://www.unef.es/es/comunicacion/comunicacion-post/el-autoconsumo-sigue-avanzando-superando-ya-el-umbral-de-los-8gw-instalados · https://www.eseficiencia.es/2026/02/02/autoconsumo-fotovoltaico-alcanza-9-3-gw-potencia-instalada-espana-segun-datos-unef
127. BOE — *Royal Decree 244/2019, of April 5* (administrative, technical and economic conditions of self-consumption; simplified surplus compensation ≤100 kW, voluntary, on the monthly bill — net billing; individual and collective self-consumption; administrative registry); MITECO — *Frequently asked questions on self-consumption*; UNEF — *Real Decreto 244/2019*; Puentes de Muras — *RD 244/2019: novedades* (extension of the collective self-consumption radius to 5 km, plants up to 5 MW). https://www.boe.es/buscar/doc.php?id=BOE-A-2019-5089 · https://www.miteco.gob.es/en/energia/energia-electrica/electricidad/autoconsumo-electrico/preguntas-frecuentes-autoconsumo.html · https://autoconsumo.unef.es/real-decreto-244-2019/ · https://puentesdemuras.com/es/rd-244-2019-normativa-fotovoltaica-espana/
128. t2p — product manual (certify any file in batches — only the fingerprint travels; screen and web-page capture with full image + source code; document scanner up to 50 pages; photo with forensic GPS; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the CAE, in chapter 23; the community minutes, in chapter 37; the invoice, in chapter 34.
129. European Parliament (EPRS) — Enrico Pastori and Raffaele Vergnani, *Odometer manipulation in motor vehicles in the EU* (2018; damage of €5,600-9,600 million/year; 30-50% in cross-border sales). https://www.europarl.europa.eu/RegData/etudes/STUD/2018/615637/EPRS_STU(2018615637_EN.pdf
130. ETSC (European Transport Safety Council) — *EU should act against mileage tampering on second-hand cars* (up to 50% of cross-border cars). https://etsc.eu/eu-should-act-against-mileage-tampering-on-second-hand-cars/
131. carVertical — *The cost of odometer fraud in Europe* (2024; ~€5.3 billion/year; average overprice of 26.3%). https://www.carvertical.com/en/blog/cost-of-odometer-fraud-in-europe
132. Car-Pass (Belgium) — *Annual report 2022* (1,344 cases in 2022, <0.2%; imports 2.4%→0.87%). https://www.car-pass.be/en/blog/car-pass-annual-report-2022

133. MITECO — *CAE measure sheet catalog: transport sector* (sheets TRA010 to TRA050). https://www.miteco.gob.es/en/energia/eficiencia/cae/catalogo-de-fichas/catalogo-vigente-de-fichas/transporte.html
134. ICCT — *Real-world use cases for zero-emission trucks: heavy tractor-trailers for goods transport in the EU* (August 2025; 65% less energy; real range +11-19%). https://theicct.org/publication/rw-use-cases-for-zet-heavy-tractor-trailers-for-goods-transport-in-the-eu-aug25/
135. Upstream Security — *2026 Global Automotive & Smart Mobility Cybersecurity Report* (494 cyberattacks in 2025; 92% remote; 67% on telematics/cloud). https://upstream.auto/press-releases/upstream-security-research-ransomware-attacks-on-automotive-and-smart-mobility-more-than-doubled-in-2025/
136. Eurojust — *Suspects arrested in Romania, Belgium and Austria for defrauding EUR 420,000 using cloned fuel cards* (2022). https://www.eurojust.europa.eu/news/suspects-arrested-romania-belgium-and-austria-defrauding-eur-420000-using-cloned-fuel-cards
137. MOBI (Mobility Open Blockchain Initiative) — *Vehicle Identity (VID) standard* (2019) and *Global Battery Passport MVP* (2024). https://dlt.mobi/mobi-announces-the-first-vehicle-identity-vid-standard-on-blockchain-in-collaboration-with-groupe-renault-ford-and-bmw-among-others/
138. Regulation (EU) 2023/1542 — *digital Battery Passport* (mandatory from 18-Feb-2027 for EV and industrial batteries above 2 kWh; QR access). https://www.circularise.com/blogs/eu-battery-passport-regulation-requirements
139. Smart Freight Centre — *GLEC Framework v3.0 / ISO 14083:2023* (freight transport emissions accounting). https://www.smartfreightcentre.org/en/about-sfc/news/a-solid-foundation-to-further-accelerate-freight-decarbonization-smart-freight-centre-releases-updated-glec-framework-version-30/
140. Regulation (EU) 910/2014 (eIDAS), Articles 41-42 — legal presumption of accuracy of the qualified timestamp. https://eur-lex.europa.eu/eli/reg/2014/910/oj
141. Susan Shaheen and Elliot Martin (UC Berkeley, Transportation Sustainability Research Center) — *Impacts of car2go: One-Way Carsharing in North America* (2016; each shared vehicle removes 7-11 private cars; VMT and GHG down). https://phys.org/news/2016-07-car-mobility-decreases-greenhouse-gas.html
142. *Does free-floating carsharing reduce private vehicle ownership? The case of SHARE NOW in European cities* (Transportation Research Part A, 2020; average of 3.3 cars replaced; range 2.1-5.3). https://pmc.ncbi.nlm.nih.gov/articles/PMC7544607/
143. BlaBlaCar — *Sustainability Report 2024* (2.5 Mt CO_2 avoided in 2024; 92 million passengers). https://newsroom.blablacar.com/news/sustainability-report-2024
144. NACTO — *A micromobility record: 157 million trips on bike and scooter share in 2023*. https://nacto.org/latest/a-micromobility-record-157-million-trips-on-bike-and-scooter-share-in-2023/
145. Insurance Journal (EDNY prosecutors) — *Two charged in $40M ride-hail surge-pricing GPS spoofing scheme* (the "Fake GPS" case, August 2024). https://www.insurancejournal.com/news/east/2024/08/30/790683.htm
146. Hai Yang and Xiaolei Wang — *Managing network mobility with tradable credits* (Transportation Research Part B, 45(3), 2011, pp. 580-594). https://ideas.repec.org/a/eee/transb/v45y2011i3p580-594.html
147. Oded Cats et al. — *Design and classification of tradable mobility credit schemes* (Transport Policy, 136, 2023) and *Tradable mobility credits for long-distance travel in Europe* (Transportation Research Part A, 186, 2024). https://ideas.repec.org/a/eee/trapol/v136y2023icp59-69.html
148. Aljamal et al. — *Tradable credit scheme for mobility management: state of the art* (Transportation Research Interdisciplinary Perspectives, 2025; no real-scale implementation to date). https://www.sciencedirect.com/science/article/pii/S2590198225002908
149. Next Mobility / ONSS — *Mobility budget in Belgium: ONSS figures 2024* (18,386 workers with a mobility budget; 17,030 without a company car). https://www.nextmobility.be/en/post/mobility-budget-belgium-onss-figures-2024/
150. Bruxelles Environnement — *Prime Bruxell'Air* (aid of €505-1,010/year for deregistering a license plate, in force in 2024). https://environnement.brussels/citoyen/services-et-demandes/primes-et-aides-financieres/la-prime-mobilite-bruxellair-pour-se-deplacer-autrement-dans-la-capitale
151. SCS Global Services / Turo — *Turo becomes first global carbon-neutral certified car-sharing marketplace* (2023). https://www.prnewswire.com/news-releases/turo-becomes-first-global-carbon-neutral-certified-car-sharing-marketplace-301799877.html

152. MOBI (Mobility Open Blockchain Initiative) — *Vehicle Identity (VID) standard on blockchain*. https://www.traffictechnologytoday.com/news/smart-cities/mobi-creates-the-first-vehicle-identity-standard-on-blockchain.html
153. Newtral — *Cómo cambió Renfe las indemnizaciones por retraso del AVE* (1-7-2024: compensation at 15 and 30 minutes eliminated for AVE, Avlo, Alvia, Euromed and Intercity; new regime 50% from 60 min and 100% from 90; 200% refund in points or 150% in vouchers; justification: liberalization and legal alignment); Renfe — *Compromiso de puntualidad* and press note on the update; Xataka — *the debate over compensating from 15 minutes* (open regulatory situation). https://www.newtral.es/renfe-indemnizacion-retraso-ave-gobiernos-pp-psoe/20250801/ · https://www.renfe.com/es/es/ayuda/compromiso-puntualidad · https://www.renfe.com/es/es/grupo-renfe/comunicacion/renfe-al-dia/sala-de-prensa/renfe-actualiza-compromiso-voluntario-puntualidad-ave-avlo-alvia-euromed-intercity · https://www.xataka.com/movilidad/renfe-esta-obligada-a-indemnizar-retrasos-15-minutos-a-partir-1-enero-gobierno-quiere-impedirlo
154. EUR-Lex — *Regulation (EC) 261/2004* (compensation of €250/400/600 for cancellation and, under the CJEU's Sturgeon doctrine, long delay ≥3 hours at arrival; extraordinary circumstances exception; right to care; CJEU case law: the "arrival time" is the opening of the doors). https://eur-lex.europa.eu/legal-content/ES/TXT/?uri=CELEX:32004R0261
155. t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; background certification that resumes when coverage returns — airplane mode; certified screen capture; document scanner; projects to group a trip's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the services claim, in chapter 33; the provider's clock, in chapter 59.
156. Coalition Against Insurance Fraud — *The Impact of Insurance Fraud on the U.S. Economy* (2022; $308.6 billion/year, ~$933 per person). https://insurancefraud.org/wp-content/uploads/The-Impact-of-Insurance-Fraud-on-the-U.S.-Economy-Report-2022-8.26.2022.pdf
157. Insurance Europe — *Annual Report 2023-2024* (Fraud chapter) and *The impact of insurance fraud* (~10% of claims spending; 2013 estimate, ~€13 billion referring to 2017). https://www.insuranceeurope.eu/downloads/ar-2024-fraud/Fraud.pdf
158. Association of British Insurers (ABI) — *Fraudulent insurance claims continue to top £1 billion* (2024: £1.16 billion; 98,400 claims). https://www.abi.org.uk/news/news-articles/2025/11/fraudulent-insurance-claims-continue-to-top-1-billion/
159. ICEA — *El sector detectó 285.407 casos sospechosos de fraude, de los que evitó un 63,7 %* (Spain, 2024). https://segurosnews.com/news/el-sector-detecto-285-407-casos-sospechosos-de-fraude-de-los-que-evito-un-637
160. Verisk (with the National Insurance Crime Bureau) — *AI editing tools are fueling a new era of insurance fraud* (March 2026; 36% of consumers, Gen Z 55%; 99% of insurers). https://www.verisk.com/company/newsroom/ai-editing-tools-are-fueling-a-new-era-of-insurance-fraud-according-to-new-research-from-verisk/
161. The Guardian (via PetaPixel) — *Fraudsters are editing photos of damaged vehicles to claim insurance* (LV= case; +300% manipulated photos, Allianz/Zurich). https://petapixel.com/2024/05/07/fraudsters-are-editing-photos-of-damaged-vehicles-to-claim-insurance/
162. Allianz UK — *Allianz UK detects frauds totalling £93m in first half of 2025* (15,800 cases). https://www.allianz.co.uk/news-and-insight/news/allianz-uk-detects-frauds-totalling-93m-in-first-half-of-2025.html
163. Zurich Insurance — *AI poses challenges, offers tantalizing solutions to insurers fighting fraud* (metadata + error level analysis). https://www.zurich.com/insights/business/ai-poses-challenges-offers-tantalizing-solutions-to-insurers-fighting-fraud
164. Swiss Re Institute — *SONAR 2025: How deepfakes, disinformation and AI amplify insurance fraud* (+700%; recommends liveness and blockchain anchoring). https://www.swissre.com/institute/research/sonar/sonar2025/how-deepfakes-disinformation-ai-amplify-insurance-fraud.html
165. Línea Directa — *VII Barómetro del Fraude al Seguro* (8.8% of motor claims in Spain). https://www.lineadirectaaseguradora.com/-/vii-barometro-de-fraude-al-seguro
166. OECD / EUIPO — *Mapping Global Trade in Fakes 2025* (counterfeits: $467 billion in 2021, 2.3% of world imports; EU 4.7%). https://www.oecd.org/en/publications/mapping-global-trade-in-fakes-2025_94d3b29f-en.html

167. World Health Organization — *Substandard and falsified medical products* (≥1 in 10 in low/middle-income countries; ~$30.5 billion/year). https://www.who.int/news-room/fact-sheets/detail/substandard-and-falsified-medical-products
168. WHO — *Medical Product Alert No. 6/2022* (The Gambia; Maiden Pharmaceuticals syrups contaminated with diethylene glycol). https://www.who.int/news/item/05-10-2022-medical-product-alert-n-6-2022-substandard-(contaminated-paediatric-medicines
169. WHO — *Medical Product Alert No. 5/2025* (India; "Coldrif" syrup by Sresan Pharmaceuticals). https://www.who.int/news/item/13-10-2025-medical-product-alert-n-5-2025--substandard-(contaminated-oral-liquid-medicines
170. EUR-Lex — *Regulation (EU) 2023/1115* (EUDR; plot geolocation coordinates; 7 commodities; cutoff date 31-Dec-2020). https://eur-lex.europa.eu/eli/reg/2023/1115/oj/eng
171. European Commission (Access2Markets) — *Delay until December 2026 — implementation of the EUDR* (second postponement, Reg. EU 2025/2650; application 30-Dec-2026 / 30-Jun-2027). https://trade.ec.europa.eu/access-to-markets/en/news/delay-until-december-2026-and-other-developments-implementation-eudr-regulation
172. EMA — *Falsified medicines: overview* (Directive 2011/62/EU; mandatory 2D unique identifier since 9-Feb-2019). https://www.ema.europa.eu/en/human-regulatory-overview/public-health-threats/falsified-medicines-overview
173. Aura Blockchain Consortium — *About* (founded in 2021 by LVMH, Prada and Cartier; >50 million products registered). https://auraconsortium.com/about
174. WWD — *Aura Blockchain crosses 50 million products* (LVMH, Prada, OTB; September 2024). https://wwd.com/business-news/technology/aura-blockchain-crosses-50-million-products-lvmh-prada-otb-1236606994/
175. Federal Trade Commission — *FTC Announces Final Rule Banning Fake Reviews and Testimonials* (August 14, 2024; in force October 21, 2024; up to $51,744 per violation; bans the sale of reviews, insider reviews without disclosure, controlled review sites, selective suppression and fake social-influence indicators); Federal Register — *Trade Regulation Rule on the Use of Consumer Reviews and Testimonials*. https://www.ftc.gov/news-events/news/press-releases/2024/08/federal-trade-commission-announces-final-rule-banning-fake-reviews-testimonials · https://www.federalregister.gov/documents/2024/08/22/2024-18519/trade-regulation-rule-on-the-use-of-consumer-reviews-and-testimonials
176. Infobae — *Amazon demanda a miles de grupos de Facebook por publicar reseñas falsas* (July 19, 2022; more than 10,000 groups; US, UK, Germany, France, Italy, Spain and Japan); Marketing4eCommerce — *Amazon denuncia a los administradores de más de 10.000 grupos de Facebook* ("Amazon Product Review," more than 43,000 members; $10 per review; more than 10,000 groups reported to Meta since 2020, half removed); La Nación — *Amazon demanda a los administradores de 10.000 grupos de Facebook*. https://www.infobae.com/america/agencias/2022/07/19/amazon-demanda-a-miles-de-grupos-de-facebook-por-publicar-resenas-falsas/ · https://marketing4ecommerce.net/amazon-reviews-falsas/ · https://www.lanacion.com.ar/tecnologia/amazon-demanda-a-los-administradores-de-10000-grupos-de-facebook-por-promover-resenas-falsas-nid19072022/
177. OCU — *Black Friday: descuentos que no son reales* (only 1 in 3 products gets cheaper; average rises of 13% in the preceding weeks in categories like smartwatches, cameras and headphones; 43% of products spike in price; advertised saving of 23% vs real prices 3% dearer; in 99% of cases the "previous price" is not the lowest of the last 30 days; monitoring from September); Público — *La trampa del Black Friday: precios inflados, rebajas ficticias*. https://www.ocu.org/consumo-familia/compras-online/noticias/bf-precios-rebajados · https://www.publico.es/sociedad/trampa-black-friday-precios-inflados-rebajas-ficticias-fiebre-compras-detiene.html
178. European Consumer Centre Spain (Ministry of Consumer Affairs) — *Spain extends the minimum legal warranty from two to three years* (since January 1, 2022; transposition of Directives (EU) 2019/770 and 2019/771; presumption of origin of the defect extended from 6 months to 2 years — 1 year for second-hand; spare parts for 10 years; limitation period from 3 to 5 years); OCU — *Garantía de los productos: se amplía el plazo*. https://portal-cec.consumo.gob.es/en/comunicacion/noticias/2022/espana-amplia-la-garantia-legal-minima-de-dos-tres-anos · https://www.ocu.org/consumo-familia/derechos-consumidor/noticias/novedades-garantia
179. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; attestation without

identification — the proof of purchase travels, the identity never; content on the device and voluntary disclosure). See chapters 13 and 15; the product's origin, in chapter 29.

180. Digital Perito — *Estafas en Wallapop y Vinted: cómo denunciar con pruebas* (>19 million active Wallapop users in Spain; second-hand market >€13 billion a year; second-hand platform scams as the bulk of the >412,000 computer frauds of 2024; +60% complaints in 2025); Público — *Así operan los estafadores en plataformas de segunda mano* (typologies: off-platform payment, triangulation/"switcheroo" exploiting buyer protection, phishing, fraudulent labels, impersonation); OKDiario / Genbeta — Guardia Civil alerts and detection tips. https://digitalperito.es/blog/estafas-wallapop-vinted-como-denunciar-pruebas-2026/ · https://www.publico.es/sociedad/asi-operan-estafadores-plataformas-segunda-mano-vinted-wallapop.html · https://okdiario.com/curiosidades/fraude-que-arrasa-plataformas-segunda-mano-guardia-civil-lanza-advertencia-sobre-vinted-wallapop-15912396 · https://www.genbeta.com/seguridad/estas-estafas-wallapop-comunes-varios-trucos-para-pillarlas-rapidamente
181. Civil framework of sales between individuals (hidden defects; absence of consumer commercial warranty). See chapter 30 for the B2C warranty and chapter 36 for the fake listing scam.
182. t2p — product manual (video with overlaid date/time; photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; certified screen capture; certify any file; projects to group a sale's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the tricked payment, in chapter 69.
183. Directorate-General for Consumer Affairs (DSCA) — *Prácticas irregulares de compañías españolas de alquiler de coches* (disproportionate indemnities of €300 to €1,000 for small scratches or blemishes, charged to the deposit card with the amount set unilaterally); CEC-España (European Consumer Centre) — *Alquiler de vehículos: derechos y recomendaciones* (later charges for cleaning or minor damage; inspections without the consumer present; recommendations: return with the office open, check with staff, request a receipt and "take photographs of the interior and the exterior as evidence"); AEA — *Líder europeo del rent-a-car condenado a devolver el importe de unos daños cobrados indebidamente*. https://www.dsca.gob.es/sites/default/files/consumo_masinfo/PRACTICAS_IRREGULARES_ALQUILER_COCHES · https://portal-cec.consumo.gob.es/sites/default/files/documentos/Folleto_Alquiler_Coches.pdf · https://portal-cec.consumo.gob.es/en/comunicacion/noticias/2021/cec-espana-recuerda-los-derechos-de-las-personas-consumidoras-cuando · https://aeaclub.org/rent-a-car-reclamacion-coche-alquiler-cobro-indebido/
184. t2p — product manual (video with overlaid date/time; photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; document scanner; projects to group a rental's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the flat's inventory, in chapter 36; the sold item's condition, in chapter 31; the card charge, in chapter 69.
185. BOE — *Ley 10/2025, de 26 de diciembre, por la que se regulan los servicios de atención a la clientela* (BOE 27-12-2025; 12 months to adapt; 95% of calls in <3 minutes; identifying tracking code; response in 15 days — 5 for unauthorized charges; 15 days' notice before automatic renewals; simple cancellation with ID and contract number; ban on exclusively automated service); La Moncloa — *Ley de Servicios de Atención a la Clientela: objetivos y medidas*; Sepín — *Novedades de la Ley 10/2025*. https://www.boe.es/buscar/act.php?id=BOE-A-2025-26698 · https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/derechos-sociales-consumo-agenda-2030/paginas/2025/ley-servicios-atencion-al-cliente.aspx · https://blog.sepin.es/blog/2026/01/14/novedades-ley-10-2025-servicios-atencion-clientela/
186. FACUA — *Multa de 20.000 euros a Vodafone: cobró a una socia de FACUA servicios de terceros que había desconectado* (right of disconnection violated up to three times; obstruction of the inspection; €20,100 with reduction); FACUA — *Multa a Movistar por acosar a través de una gestora de cobros a una socia por una deuda falsa* (€1,741 for lines never contracted); FACUA — *El Supremo confirma la multa de 765.000 € a Telefónica por limitar el derecho a recibir facturas en papel*; Xataka Móvil — *FACUA denuncia a Telefónica por la penalización por baja anticipada*. https://facua.org/noticias/multa-de-20-000-euros-a-vodafone-cobro-a-una-socia-de-facua-servicios-de-terceros-que-habia-desconectado/ · https://www.facua.org/es/noticia.php?Id=9279 · https://facua.org/noticias/el-supremo-confirma-la-multa-de-765000-de-la-junta-de-andalucia-a-telefonica-por-limitar-el-derecho-a-recibir-facturas-en-papel/ ·

https://www.xatakamovil.com/conectividad/facua-denuncia-a-telefonica-por-la-penalizacion-por-baja-anticipada-en-movistar-fusion

187. t2p — product manual (background audio recording; certified screen capture and recording with foreground app verification; web page capture with full image + source code; document scanner up to 50 pages; projects to group captures by case; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; the lawfulness of recording one's own conversations, in chapter 75; the cloned voice, in chapter 68.
188. AEAT — *Sistemas Informáticos de Facturación (SIF) y VERI*FACTU *(RD 1007/2023: standardized invoicing record with chained fingerprint — each record contains the previous one's hash — electronic signature and verification QR; "Verifactu" immediate voluntary submission mode and unalterable preservation; official FAQ); Guía Fiscal* — Verifactu 2026-2027 *and Davisa* — Guía Verifactu 2026* (calendar postponed by RD-Law 15/2025, of December 2: 1-1-2027 companies, 1-7-2027 the rest; penalties of up to €50,000/year for users and €150,000 for manufacturers). https://sede.agenciatributaria.gob.es/Sede/iva/sistemas-informaticos-facturacion-verifactu.html · https://sede.agenciatributaria.gob.es/Sede/iva/sistemas-informaticos-facturacion-verifactu/preguntas-frecuentes/sistemas-verifactu.html · https://guiafiscal.es/autonomos/verifactu-2026-guia/ · https://www.davisa.es/guia-verifactu-2026/
189. BOE — *Ley 3/2004, de 29 de diciembre, de medidas de lucha contra la morosidad en las operaciones comerciales* (60 calendar days maximum by agreement between companies, 30 by default from receipt; automatic late-payment interest; €40 collection costs per invoice; public administrations 30 days — Law 15/2010); Togas — *Los intereses de demora en el primer semestre de 2025: 11,15 %* (Treasury Resolution of 30-12-2024); PMcM / Revista Limpiezas — *2025 late-payment report* (85% of large companies breach the deadlines; public administrations at 70 days on average, +3 in 2025; private sector at 67). https://www.boe.es/buscar/act.php?id=BOE-A-2004-21830 · https://togas.biz/articulos/articulo-profesionales--los-intereses-de-demora-en-las-operaciones-comerciales-en-el-primer-semestre-de-2025-quedan-fijados-en-un-11-15-por-resolucion-de-30-de-diciembre-de-2024-de-la-secretar-a-general-del-tesoro-y-financiacion-internacional/ · https://www.revistalimpiezas.es/actualidad/normativa-limpiezas/informe-morosidad-pmcm-el-85-de-las-grandes-companias-incumple-los-plazos-de-pago-y-se-financia-en-parte-a-costa-de-las-pymes-segun-la-pmcm_20260417.html · https://pmcm.es/informe-morosidad-estudio-plazos-pago-espana-2024/
190. t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, mocked location blocks certification; background certification that resumes when coverage returns; document scanner up to 50 pages in one PDF; certified screen and email capture; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; international trade, in chapter 54; auditing, in chapter 55.
191. World Bank — Laura Tuck and Wael Zakout, *7 reasons for land and property rights to be at the top of the global agenda* (~30% of the population with registered title). https://blogs.worldbank.org/en/voices/7-reasons-land-and-property-rights-be-top-global-agenda
192. Prindex — *Tenure insecurity on the rise* (2024; ~1.1 billion adults, 23%, insecure about their land/housing). https://www.prindex.net/news-and-stories/tenure-insecurity-on-the-rise/
193. FAO (SOFA report, via UN News) — *Global report highlights slow progress in expanding secure land tenure* (~35% of land documented; >1,100 M fear losing it). https://news.un.org/en/story/2026/02/1167037
194. FBI IC3 (via the National Association of Realtors) — *Online real estate fraud climbed to $275M in 2025* (12,368 victims). https://www.nar.realtor/magazine/real-estate-news/online-real-estate-fraud-climbed-to-275m-in-2025-fbi-says
195. American Land Title Association (ALTA) — *Seller Impersonation Fraud Attempts Increase, Study Shows* (July 2024; 28%; vacant land 85%; fake notary credentials 43%). https://www.alta.org/news-and-publications/news/20240730-Seller-Impersonation-Fraud-Attempts-Increase-Study-Shows
196. FBI Boston (via ALTA) — *FBI Boston Issues Quit Claim Deed Fraud Warning* (April 2025; real estate fraud 2019-2023: 58,141 victims / $1.3 billion). https://www.alta.org/news-and-publications/news/20250410-FBI-Boston-Issues-Quit-Claim-Deed-Fraud-Warning
197. Exonum / Bitfury — *The story of Georgia* (NAPR; ~1.5 million titles anchored to Bitcoin; "no one can alter the timestamp"). https://exonum.com/story-georgia

198. Eurasianet — *Georgia authorities use blockchain technology for developing land registry* (extension Feb-2017; Hernando de Soto, adviser). https://eurasianet.org/georgia-authorities-use-blockchain-technology-for-developing-land-registry
199. Computer Weekly — *Sweden trials blockchain for land registry management* (Lantmäteriet + ChromaWay; >€100 million/year in savings). https://www.computerweekly.com/news/450421958/Sweden-trials-blockchain-for-land-registry-management
200. CoinDesk — *Blockchain land title project stalls in Honduras* (Factom; stalled on political questions, 2015). https://www.coindesk.com/markets/2015/12/26/blockchain-land-title-project-stalls-in-honduras
201. Libertad Digital — *Cae en Cádiz una red criminal por una «estafa inmobiliaria a nivel nacional»* (May 6, 2026; Operation ISLEDA; 3 arrests aged 19-34 in San Fernando; uncovered in Valladolid; accounts opened with identities stolen in fake job offers; fraud, laundering and criminal organization); Cádiz Directo — *La Guardia Civil destapa en San Fernando una trama nacional de falsas viviendas de alquiler*. https://www.libertaddigital.com/andalucia/2026-05-06/san-fernando-cadiz-valladolid-guardia-civil-detiene-tres-hombres-acusados-de-estafa-inmobiliaria-conocido-portal-web-alquiler-y-venta-de-pisos-utilizaban-sus-datos-para-cuentas-falsas-victimas-por-todo-el-pais-destapado-en-valladolid-1b-7400307/ · https://www.cadizdirecto.com/cadiz/la-guardia-civil-destapa-en-san-fernando-una-trama-nacional-de-falsas-viviendas-de-alquiler
202. Idealista/News — *La última estafa en los anuncios de alquiler y los consejos para no caer en ella* (bait prices, owners "abroad," payments before the visit); Maldita.es — *Anuncios de alquiler falsos e inmobiliarias fantasma: cómo reconocer timos inmobiliarios*. https://www.idealista.com/news/inmobiliario/vivienda/2022/11/18/800160-la-ultima-estafa-en-los-anuncios-de-alquiler-y-los-consejos-para-no-caer-en-ella · https://maldita.es/malditobulo/20230830/inmobiliaria-fantasma-alquiler-piso-fraude-verano/
203. Idealista/News — *Todo lo que tienes que saber sobre la devolución de la fianza del alquiler* (one-month deadline; statutory interest; settlement agreement with inventory); Ático Jurídico — *Devolución de la fianza del alquiler* (courts demand proof from whoever withholds: invoices, reports); Fotocasa Life — *Cómo proceder si no nos devuelven la fianza del alquiler*. https://www.idealista.com/news/inmobiliario/vivienda/2022/07/03/797535-todo-lo-que-tienes-que-saber-sobre-la-devolucion-de-la-fianza-del-alquiler · https://aticojuridico.com/devolucion-fianza-alquiler/ · https://www.fotocasa.es/fotocasa-life/alquiler/como-proceder-si-no-nos-devuelven-la-fianza/
204. Housage — *Ley antiokupas 2025: desalojo en 15 días, cómo funciona y a quién no protege* (LO 1/2025 in force April 3, 2025; breaking-and-entering Art. 202 CP and usurpation Art. 245 CP via fast-track trial; police eviction within the first 48 hours; abolition of the vulnerability report on this track); Blegal — *La Ley 1/2025 agiliza el desalojo de okupas*; Vento Abogados — *Juicio rápido por okupación: la nueva regulación de la LO 1/2025*. https://housage.es/articulos/actualidad-inmobiliaria/ley-antiokupas-2025-desalojo-en-15-dias-como-funciona-y-a-quien-no-protege/ · https://www.blegal.eu/la-reforma-de-la-ley-organica-1-2025-agilizara-los-juicios-rapidos-para-facilitar-el-desalojo-de-okupas/ · https://vento.es/juicio-rapido-por-okupacion/
205. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; content never leaves the device; voluntary disclosure; the gaps in the series are dated too). See chapters 13 and 15; for title and the registry, chapter 35.
206. Vecindia — *Condenado administrador de fincas por apropiación indebida* (Gijón: 1 year and a half in prison upheld; €2,984.04 from the community account; deceptive maneuvers in the minute book getting former presidents to sign; fake works and locksmith receipts as cover; discovered when the president and treasurer changed); ADADE — *El Tribunal Supremo ratifica la condena por apropiación indebida a un administrador de fincas*; El Derecho — *Responsabilidad penal del administrador de fincas* (Art. 253 CP). https://vecindia.es/comunidades-de-propietarios/condenado-administrador-fincas-apropiacion-indebida/ · https://www.adade.es/comunicacion/sala-de-prensa/el-tribunal-supremo-ratifica-la-condena-por-apropiacion-indebida-a-un-administrador-de-fincas · https://elderecho.com/cuando-puede-concurrir-responsabilidad-penal-en-el-administrador-de-fincas-por-su-actuacion-en-una-comunidad-de-propietarios/
207. Horizontal Property Act (challenge of agreements that are unlawful, harmful or abusive; notices and minutes; remote meetings enabled by recent reform); Mundojurídico — *Responsabilidad del administrador de la Comunidad*; Oliva-Ayala — *Falsedad documental y certificaciones de juntas*.

https://mundojuridico.info/responsabilidad-del-administrador-de-la-comunidad/ · https://www.oliva-ayala.es/wp-content/uploads/2024/03/Alcacer-2024-LLPenal-Falsedades.pdf
208. t2p — product manual (document scanner up to 50 pages in one certified PDF; audio recording; certified screen and email capture; photo with forensic GPS; projects to group a works file; certify any file — only the fingerprint travels; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; property and the registry, in chapter 35; renting, in chapter 36.
209. Atlántico — *Investigan varias estafas en reformas de viviendas en Vigo y su comarca* (at least six victims; up to €10,000 each; a builder who starts works and disappears); Liberal de Castilla — *Un vecino de Guadalajara pierde sus ahorros ante un falso albañil de internet* (Guardia Civil; scammer located in Madrid); Telemadrid — *Destapamos al presunto estafador de las reformas*; Allende Abogados — *Estafas en reformas* (pattern: attractive digital ad, cheap quote, advance "for materials," disappearance; losses exceeding €77,000 in some cases). https://www.atlantico.net/articulo/vigo/investigan-estafa-reformas-menos-media-docena-victimas/20221017222637946245.html · https://www.liberaldecastilla.com/un-vecino-de-guadalajara-pierde-sus-ahorros-ante-un-falso-albanil-de-internet-la-guardia-civil-acorrala-al-estafador-en-madrid/ · https://www.telemadrid.es/programas/la-redaccion/Destapamos-presunto-estafador-reformas-2-2339786043--20210510052540.html · https://www.allendeabogados.com/estafas-reformas/
210. Framework of the works contract (Civil Code) and consumer law applicable to renovations for consumers (prior quote and invoice). See chapter 34 for the professional's collection and chapter 37 for communal works.
211. t2p — product manual (document scanner up to 50 pages in one certified PDF; photo with forensic GPS — genuine fix ~50 m / ≤30 s; video with overlaid date/time; projects to link quote, phases, evidence and payments; signature with video selfie, liveness challenges and biometrics; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the levy, in chapter 37; the invoice, in chapter 34.
212. BCG and ADDX — *Relevance of On-chain Asset Tokenization in 'Crypto Winter'* (2022; ~$16.1 trillion by 2030, ~10% of world GDP). https://www.bcg.com/publications/2022/relevance-of-on-chain-asset-tokenization
213. McKinsey — *From ripples to waves: the transformational power of tokenizing assets* (2024; base case ~$2 trillion, range 1-4, excluding stablecoins). https://www.mckinsey.com/industries/financial-services/our-insights/from-ripples-to-waves-the-transformational-power-of-tokenizing-assets
214. CoinDesk (RedStone / Gauntlet / RWA.xyz report) — *Real-world asset tokenization market has grown almost fivefold in 3 years* (~$24 billion; private credit ~$14 billion). https://www.coindesk.com/business/2025/06/26/real-world-asset-tokenization-market-has-grown-almost-fivefold-in-3-years
215. BlackRock / BusinessWire — *BlackRock Launches Its First Tokenized Fund, BUIDL, on the Ethereum Network* (20-Mar-2024). https://www.businesswire.com/news/home/20240320771318/en/BlackRock-Launches-Its-First-Tokenized-Fund-BUIDL-on-the-Ethereum-Network
216. Securitize / PR Newswire — *BlackRock's BUIDL surpasses $1B in AUM* (March 2025). https://www.prnewswire.com/news-releases/blackrock-usd-institutional-digital-liquidity-fund-buidl-tokenized-by-securitize-surpasses-1b-in-aum-302401480.html
217. Paxos — *PAXG transparency* (serialized bars at Brink's London; monthly KPMG attestations). https://www.paxos.com/paxg-transparency
218. Financial Stability Board — *The Financial Stability Implications of Tokenisation* (October 2024; custodians and oracles as new points of vulnerability). https://www.fsb.org/2024/10/the-financial-stability-implications-of-tokenisation/
219. Moody's (via CoinDesk) — *Tokenized fund adoption grows but brings technology risks* (2024; oracle failures, limited track record). https://www.coindesk.com/markets/2024/01/15/tokenized-fund-adoption-grows-but-brings-technology-risks-moodys
220. Standard Chartered and Synpulse — *Trade finance to play substantial role in $30.1 trillion tokenised real-world assets market by 2034*. https://www.sc.com/en/press-release/trade-finance-to-play-substantial-role-in-usd-30-1-trillion-tokenised-real-world-assets-market-by-2034/

221. InvestmentNews — *Larry Fink takes aim at investment inequality* (BlackRock's 2025 annual letter; identity verification as the condition). https://www.investmentnews.com/alternatives/larry-fink-takes-aim-at-investment-inequality/259910
222. NPR — *The New York Times sues OpenAI and Microsoft for copyright infringement* (27-Dec-2023; 127-page exhibit with GPT-4 outputs). https://www.npr.org/2023/12/27/1221821750/new-york-times-sues-chatgpt-openai-microsoft-for-copyright-infringement
223. Bloomberg Law — *OpenAI must turn over 20 million ChatGPT logs, judge affirms* (January 2026). https://news.bloomberglaw.com/ip-law/openai-must-turn-over-20-million-chatgpt-logs-judge-affirms
224. Getty Images — *Getty Images statement* (action against Stability AI; consent and licensing, January 2023). https://newsroom.gettyimages.com/en/getty-images/getty-images-statement
225. Judiciary of England and Wales — *Getty Images v Stability AI [2025] EWHC 2863 (Ch)* (Nov 2025; dismisses most of the copyright claims). https://www.judiciary.uk/judgments/getty-images-v-stability-ai/
226. Copyright Alliance — *Andersen v. Stability AI* (Sarah Andersen, Kelly McKernan, Karla Ortiz; discovery since August 2024). https://copyrightalliance.org/andersen-v-stability-ai-copyright-case/
227. Karla Ortiz — *Testimony before the U.S. Senate Judiciary Committee* (12-Jul-2023; "never asked, never credited, never compensated"). https://www.judiciary.senate.gov/imo/media/doc/2023-07-12_pm_-testimony-_ortiz.pdf
228. Variety — *AI-generated fake Drake/The Weeknd song 'Heart on My Sleeve'* (April 2023; pulled by Universal Music Group). https://variety.com/2023/music/news/fake-ai-generated-drake-weeknd-collaboration-heart-on-my-sleeve-1235585451/
229. Wikipedia — *2023 SAG-AFTRA strike* (14-Jul to 9-Nov-2023; consent for the "digital replica"). https://en.wikipedia.org/wiki/2023_SAG-AFTRA_strike
230. U.S. Copyright Office — *Copyright and Artificial Intelligence, Part 2: Copyrightability* (29-Jan-2025; human authorship as the foundation). https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf
231. USPTO — *America Invents Act (AIA): Frequently Asked Questions* (first to file since 16-Mar-2013; "the date of invention is irrelevant"). https://www.uspto.gov/patents/laws/america-invents-act-aia/america-invents-act-aia-frequently-asked
232. WIPO (World Intellectual Property Organization) — *Patents: Frequently Asked Questions* (prior art; prior disclosure destroys novelty; the grace period is not universal). https://www.wipo.int/en/web/patents/faq_patents
233. BitLaw — *What does a patent application cost?* (drafting $9,000-17,000; budget $15,000-25,000 before grant). https://www.bitlaw.com/guidance/patent/what-does-a-patent-application-cost.html
234. USPTO — *Fee Schedule* (provisional application: fixes the date, 12 months; official fees). https://www.uspto.gov/sites/default/files/documents/USPTO-fee-schedule_current.pdf
235. McLane Middleton — *Know the Law: Is the 'Poor Man's Patent' Real?* (no legal value; obsolete since the 2013 first-to-file). https://www.mclane.com/insights/know-the-law-is-the-poor-mans-patent-real/
236. Technical Disclosure Commons — *FAQ* (free; automatic publication date; indexed by Google Patents). https://www.tdcommons.org/faq.html
237. IP.com — *Publishing Disclosures to the Prior Art Database* (timestamping via Surety AbsoluteProof). https://kb.ip.com/pad/knowledge-base/publishing-disclosures-to-the-prior-art-database-pad/
238. Richard Coller III and Alexander Covington (IPWatchdog) — *Defensive Publications: A Cost-Effective Tool to Supplement Your Patent Strategy* (2020). https://ipwatchdog.com/2020/05/25/defensive-publications-cost-effective-tool-supplement-patent-strategy/
239. European Patent Office (EPO) — *Espacenet reaches 160 million patent documents* (April 2025). https://agip-news.com/news.aspx?id=76074&lang=en
240. WIPO — *PATENTSCOPE* (~128 million patent documents; 5.4 million PCT applications). https://patentscope2.wipo.int/
241. CNBC — *Trump pardons Anthony Levandowski, the engineer who stole self-driving car secrets from Google* (thousands of files downloaded before leaving; Otto bought by Uber; 2018 Waymo-Uber settlement valued by the press at ~$245 million in shares; guilty plea in March 2020; 18 months in prison — Judge William Alsup: "the biggest trade secret crime I have ever seen"; full pardon on 20-1-2021, the last day of the term); TechCrunch — *Former Google engineer Anthony Levandowski among list of last-minute Trump pardons*; Wikipedia — *Anthony Levandowski* (2022

global settlement: $25-30 million). https://www.cnbc.com/2021/01/20/anthony-levandowski-pardoned-after-stealing-trade-secrets-from-google.html · https://techcrunch.com/2021/01/19/former-google-engineer-anthony-levandowski-among-list-of-last-minute-trump-pardons/ · https://en.wikipedia.org/wiki/Anthony_Levandowski

242. BOE — *Ley 1/2019, de 20 de febrero, de Secretos Empresariales* (transposition of Directive (EU) 2016/943; three requirements: secrecy + business value + reasonable measures; scope: formulas, processes, databases, client lists, algorithms, know-how; balance: it does not cover the experience and skills honestly acquired by the worker and cannot restrict mobility; procedural confidentiality measures); Garrigues — *Aprobada la ley que protege los secretos empresariales*. https://www.boe.es/buscar/act.php?id=BOE-A-2019-2364 · https://www.garrigues.com/en_GB/new/law-will-protect-trade-secrets-has-been-approved
243. t2p — product manual (certify any file — only the fingerprint travels, the content never leaves the device; document scanner up to 50 pages; signature with video selfie, liveness challenges and biometrics; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package with selective disclosure verifiable offline). See chapters 13 and 15; prior art and the patent, in chapter 41; an attack's logs, in chapter 57; whistleblowers, in chapter 62.
244. Spend Network / World Bank — *The $13 trillion global value of public procurement* (~13-15% of GDP; under 3% in open data). https://www.spendnetwork.com/13-trillion-the-global-value-of-public-procurement/
245. World Bank — *The hidden $1 trillion: halting waste in public procurement* (~a quarter wasted; ~$1 trillion/year). https://blogs.worldbank.org/en/voices/hidden-1-trillion-halting-waste-public-procurement
246. CoST – Infrastructure Transparency Initiative — *This is why construction is so corrupt* (10-30% of public works value; ~$6 trillion/year toward 2030). https://infrastructuretransparency.org/2016/02/05/this-is-why-construction-is-so-corrupt/
247. OECD — *Guidelines for Fighting Bid Rigging in Public Procurement (2025 Update)* (collusion can raise contract prices 20% or more). https://www.oecd.org/en/publications/2025/09/oecd-guidelines-for-fighting-bid-rigging-in-public-procurement-2025-update_127880ea.html
248. Transparency International — *Corruption Perceptions Index 2025* (world average 42/100, all-time low; Denmark 89, eighth year at the top). https://www.transparency.org/en/news/cpi-2025-findings-insights-corruption
249. Open Contracting Partnership — *The impact of ProZorro in Ukraine* (~$6 billion saved; Dozorro, 33,000+ tenders). https://www.open-contracting.org/impact-stories/impact-ukraine/
250. e-Estonia — *KSI Blockchain* (first state with blockchain in production, 2012; Guardtime; only the fingerprint is anchored). https://e-estonia.com/solutions/cyber-security/ksi-blockchain/
251. World Economic Forum — *Exploring Blockchain Technology for Government Transparency to Reduce Corruption* (Colombia, School Feeding Program, IDB, 2020). https://www.weforum.org/publications/exploring-blockchain-technology-for-government-transparency-to-reduce-corruption/
252. EUR-Lex — *eIDAS Regulation (EU 910/2014), Article 41* (legal presumption of accuracy of the qualified timestamp). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2014.257.01.0073.01.ENG
253. World Bank — *GovTech Maturity Index* (198 economies; four areas of digital government maturity). https://www.worldbank.org/en/programs/govtech/gtmi
254. BOE — *Ley 39/2015, de 1 de octubre, del Procedimiento Administrativo Común de las Administraciones Públicas* (Art. 21: general three-month term to decide and notify; administrative silence with effects from expiry; silence certificate issued of the body's own motion within 15 days; electronic notifications: 10 calendar days from availability, after which deemed rejected); Administrativando — *Silencio administrativo: qué es y cuáles son sus efectos*; Conceptos Jurídicos — *Silencio administrativo en España*. https://www.boe.es/buscar/pdf/2015/BOE-A-2015-10565-consolidado.pdf · https://administrativando.es/silencio-administrativo/ · https://www.conceptosjuridicos.com/silencio-administrativo/
255. Noticias Jurídicas — *El mercado negro de citas para extranjería: venta de citas para asilo internacional* (bots hoarding appointments; resale; system overwhelmed); OKDiario — *El caos en la web de Extranjería impulsa a las mafias que cobran hasta 350 € por cita* (WhatsApp resale from €50 to €350, up to €600); Parainmigrantes — *El problema de la cita previa llega al Defensor del Pueblo* (constant complaints, ≈3,000 in a year; >163,000 asylum applications in 2023; "excessive delays"); Iberoeconomía — *España perseguirá a las mafias que revenden citas para trámites*

públicos. https://noticias.juridicas.com/contenido-patrocinado/articulo/19314-el-mercado-negro-de-citas-para-extranjeria:-venta-de-citas-para-asilo-internacional/ · https://okdiario.com/espana/caos-web-extranjeria-impulsa-mafias-que-cobran-300-cita-cada-inmigrante-13835137 · https://www.parainmigrantes.info/el-problema-de-la-gestion-de-la-cita-previa-para-tramites-de-extranjeria-llega-al-defensor-del-pueblo/ · https://iberoeconomia.es/inmigracion/espana-perseguira-a-las-mafias-que-revende-citas-para-tramites-publicos/

256. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; content in each party's systems, only fingerprints on the chain; trail verifiable by both parties). See chapters 13 and 15; public procurement, in chapter 43.
257. AFAPNA — *Oposiciones 2026: sorteo de letras de tribunales y aspirantes*; OpositaTest — *Sorteo por letra del orden de actuación en oposiciones* (public letter draw; regulated subsequent procedures; "mirror" system inverting the ID order); USTEA — *Resultado del sorteo para la designación de vocales de los tribunales*. https://afapna.es/educa/oposiciones-2026-sorteo-letras-de-tribunales-y-aspirantes/ · https://blog.opositatest.com/dudas-letra-del-orden-de-actuacion-de-los-opositores/ · https://educacion.ustea.org/oposiciones-2026-resultado-del-sorteo-para-la-designacion-de-vocales-de-los-tribunales/
258. esPúblico — *La notificación en el procedimiento de adjudicación de viviendas protegidas mediante sorteo* (draw before a notary under each administration's procedure; possibility of challenge); elDiario.es (Canarias Ahora) — *Los Llanos espera que la impugnación del sorteo de las 40 viviendas se resuelva «con prontitud y máxima transparencia»*. https://www.administracionpublica.com/adjudicaciondeviviendasprotegidas/ · https://www.eldiario.es/canariasahora/lapalmaahora/municipios/los_llanos_de_aridane/los-llanos-espera-impugnacion-sorteo-viviendas-resuelva-prontitud-maxima-transparencia_1_4583333.html
259. t2p — product manual (SHA-512 fingerprint; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; prior commitment by fingerprint — personal data are not exposed; anchored, reverifiable records). See chapters 13 and 15; electoral integrity, in chapter 48.
260. Cloudflare / League of Entropy — *Distributed Randomness Beacon* (drand: distributed, publicly verifiable randomness beacon; founding members: Cloudflare, EPFL, Kudelski Security, Protocol Labs and the University of Chile; unbiased, unpredictable, verifiable values at fixed intervals; threshold cryptography; applications: lotteries and elections); PR Newswire — *The League of Entropy launches production drand network* (the first publicly verifiable beacon in production); GitHub — *drand/drand*. https://www.cloudflare.com/leagueofentropy/ · https://www.prnewswire.com/news-releases/the-league-of-entropy-launches-production-drand-network-providing-the-first-publicly-verifiable-distributed-randomness-beacon-301109096.html · https://github.com/drand/drand
261. Redacción Médica — *Un informe revela filtraciones de 19 exámenes en hasta 12 OPE de Osakidetza* (2016-2017 public employment offer of the Basque health service; union complaints in Anesthesia and Angiology; signs in 19 exams across 12 specialties; resignations of health minister Jon Darpón, Osakidetza's director general and two HR executives); Deia — *Archivan la causa por las presuntas filtraciones en la OPE de Osakidetza en 2018 por falta de pruebas* (December 2024: shelved "as no rational indications of the commission of the offense exist," provisional for two suspects from Plastic Surgery); elDiario.es — *La Justicia cierra sin culpables la investigación* and *Ocho años del inicio de la investigación... sin fecha para el juicio*. https://www.redaccionmedica.com/autonomias/pais-vasco/un-informe-revela-filtraciones-de-19-examenes-en-hasta-12-ope-de-osakidetza-4796 · https://www.deia.eus/actualidad/sociedad/2024/12/12/archivan-causa-filtraciones-ope-osakidetza-9042604.html · https://www.eldiario.es/euskadi/justicia-cierra-responsabilidad-investigacion-filtraciones-oposiciones-medicas-vascas_1_11896007.html
262. t2p — product manual (certify any file — only the fingerprint travels, the content never leaves custody; photo with a certain date; document scanner; capture of official publications with a date; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the tribunal draw, in chapter 45; credentials, in chapter 19; whistleblowers, in chapter 62.
263. World Bank (ID4D) — *Global progress in identification: 3 findings from the latest data* (2025; ~800 M without official identity; >50% are unregistered children).

https://blogs.worldbank.org/en/digital-development/global-progress-in-identification--3-findings-from-the-latest-da
264. UNHCR — *Global Trends Report 2024* (4.4 million officially stateless, up to ~10 M real; Rohingya ~1.8 M). https://www.unhcr.org/sites/default/files/2025-06/global-trends-report-2024.pdf
265. UNICEF — *Birth registration steadily increases worldwide, but 150 million children still invisible* (December 2024). https://www.unicef.org/press-releases/birth-registration-steadily-increases-worldwide-150-million-children-still-invisible
266. European Commission — *European Digital Identity (eIDAS 2.0)* (Reg. EU 2024/1183; EUDI Wallet mandatory by the end of 2026; selective disclosure). https://digital-strategy.ec.europa.eu/en/policies/eudi-regulation
267. W3C — *Verifiable Credentials Data Model v2.0 becomes a W3C Recommendation* (15-May-2025) and *Decentralized Identifiers (DID) v1.0* (2022). https://www.w3.org/press-releases/2025/verifiable-credentials-2-0/
268. Government of India (Press Information Bureau) — *Aadhaar* (~1.34 billion living holders; >99.9% of adults; >170 billion authentications). https://www.pib.gov.in/PressReleasePage.aspx?PRID=2241778
269. Business Standard / Resecurity — *Aadhaar data of millions of Indians put on sale on the dark web* (~815 million records, October 2023). https://www.business-standard.com/india-news/aadhaar-data-of-millions-of-indians-put-on-sale-on-the-dark-web-reports-123103000993_1.html
270. Scroll.in — *Denied food because she did not have an Aadhaar-linked ration card, Jharkhand girl dies of starvation* (Santoshi Kumari, 2017). https://scroll.in/article/854225/denied-food-because-she-did-not-have-aadhaar-linked-ration-card-jharkhand-girl-dies-of-starvation
271. Reetika Khera (IIT Delhi) — research on Aadhaar and exclusion in welfare programs (*Economic & Political Weekly*). https://www.epw.in/engage/article/aadhaar-failures-food-services-welfare
272. MIT Technology Review — *Inside the Jordan refugee camp that runs on blockchain* (WFP Building Blocks; EyePay iris recognition). https://www.technologyreview.com/2018/04/12/143410/inside-the-jordan-refugee-camp-that-runs-on-blockchain/
273. Edward Snowden — *Vigilancia permanente* (Planeta, 2019).
274. Biometric Update / ICLG — *Worldcoin: bans and deletion orders* (2024-2025). https://www.biometricupdate.com/202510/colombia-orders-world-shut-down-citing-biometrics-compliance-failures
275. Wikipedia (with BBC / Al Jazeera) — *2017 Kenyan presidential election petition* (4-2 annulment; CJ David Maraga; first in Africa). https://en.wikipedia.org/wiki/2017_Kenyan_presidential_election_petition
276. The Standard (Kenya) — *Presidential petition: what scrutiny of key IEBC forms revealed* (291 34B forms; 56 without security features, 31 without serial numbers, 32 unsigned). https://www.standardmedia.co.ke/article/2001253102/presidential-petition-what-scrutiny-of-key-iebc-forms-revealed
277. NYU Center on International Cooperation / Congo Research Group — *Who really won the Congolese elections?* (Fayulu ~60%; CENCO 39,824 observers; correlation 0.976-0.991; ~3.6 million votes). https://cic.nyu.edu/resources/who-really-won-the-congolese-elections/
278. Sunoo Park, Michael Specter, Neha Narula and Ronald L. Rivest — *Going from bad to worse: from Internet voting to blockchain voting* (Journal of Cybersecurity, 2021). https://academic.oup.com/cybersecurity/article/7/1/tyaa025/6137886
279. MIT News / USENIX Security — *Researchers find security flaws in Voatz; "The Ballot is Busted Before the Blockchain"* (2020). https://news.mit.edu/2020/voting-voatz-app-hack-issues-0213
280. Pierrick Gaudry (CNRS/INRIA) — *Breaking the encryption scheme of the Moscow Internet voting system* (2019; broken in ~20 minutes). https://arxiv.org/abs/1908.05127
281. Philip B. Stark (UC Berkeley) — *A Gentle Introduction to Risk-Limiting Audits* (2012). https://www.stat.berkeley.edu/~stark/Preprints/gentle12.pdf
282. Election Academy (University of Minnesota) — *Colorado successfully completes the nation's first risk-limiting audit* (2017; 9% risk limit). https://electionacademy.lib.umn.edu/2017/11/27/colorado-successfully-completes-nations-first-risk-limiting-audit/
283. OSCE/ODIHR — *Election Observation Handbook* (6th edition; short- and long-term observers; distinguishing the isolated from the systematic). https://cdn.osce.org/sites/default/files/f/documents/5/e/68439.pdf
284. The Carter Center — *Election Standards and the Declaration of Principles for International Election Observation* (2005). https://electionstandards.cartercenter.org/our-mission/

285. Shoshana Zuboff — *La era del capitalismo de la vigilancia* (Paidós, 2020).
286. L2BEAT; MEV Watch (2026). https://www.mevwatch.info/
287. WFP — *Building Blocks* (~$555 million in transfers; ~6 million people; Ukraine: $337 million coordinated, ~$35 million of overlap avoided; 159 organizations in 2025). https://www.wfp.org/building-blocks
288. UN Innovation Network — *How the Building Blocks blockchain network is transforming humanitarian aid* (>1 million refugees each month). https://www.uninnovation.network/blog/how-the-building-blocks-blockchain-network-is-transforming-humanitarian-aid
289. UNHCR — *UNHCR's Biometric Tools in 2023* (database of 15.7 million people; PRIMES/BIMS). https://www.unhcr.org/blogs/unhcrs-biometric-tools-in-2023/
290. Development Initiatives — *Cash and Voucher Assistance* (~$10.6 billion in 2022 → $9.7 billion in 2023; WFP, largest provider). https://devinit.org/resources/falling-short-humanitarian-funding-reform/cash-voucher-assistance/
291. CALP Network — *The State of the World's Cash 2023*. https://www.calpnetwork.org/collection/the-state-of-the-worlds-cash-2023-report/
292. Human Rights Watch — *UN Shared Rohingya Data Without Informed Consent* (Lama Fakih; ~830,000 names shared with Myanmar). https://www.hrw.org/news/2021/06/15/un-shared-rohingya-data-without-informed-consent
293. UN News — *Yemen: WFP pauses food distributions* (biometric dispute with the Houthis, 2019-2023). https://news.un.org/en/story/2023/12/1144417
294. GiveDirectly / The New Humanitarian — *DRC mobile-cash fraud* (~$1.2 million diverted; ~1,900 families). https://www.givedirectly.org/drc-case-2023/
295. Carmela Troncoso et al. (EPFL SPRING Lab, with the ICRC) — *Janus: Safe Biometric Deduplication for Humanitarian Aid Distribution* (IEEE S&P 2024; reveals only one bit, no biometric database). https://arxiv.org/abs/2308.02907
296. FTC (2024); Autoriteit Persoonsgegevens / EDPB (2024). https://www.edpb.europa.eu/news/national-news/2024/dutch-supervisory-authority-imposes-fine-clearview-because-illegal-data_en
297. Wikipedia (ES) — *Terremotos de Venezuela de 2026* (M7.2 San Felipe + M7.5 Yumare doublet, 24-6-2026, 39 s apart; >3,342 dead, >16,840 injured, >68,000 missing; Simón Bolívar airport suspended; distribution conflicts: supplies held back in Altamira, civilian collection centers banned in Charallave and Campo Elías, distribution blocked in El Limón, restricted access in La Guaira; response: >30 countries, >3,600 rescuers, >1,000 tonnes); ReliefWeb — *Terremoto en Venezuela: panorama de la situación (25-6-2026)*; SkyAlert — *Actualización del terremoto en Venezuela*. https://es.wikipedia.org/wiki/Terremotos_de_Venezuela_de_2026 · https://reliefweb.int/report/venezuela-bolivarian-republic/terremoto-en-venezuela-panorama-de-la-situacion-25-de-junio-de-2026 · https://skyalert.mx/articulos/actualizacion-sismos-venezuela-2026
298. Ayuda en Acción — *Terremotos en Venezuela: qué ha pasado y situación actual* (humanitarian context; donations: José Andrés/World Central Kitchen $1 million, Rosalía $2 million to UNICEF, Metallica $100,000). https://ayudaenaccion.org/blog/ayuda-humanitaria/terremotos-venezuela/
299. U.S. Department of the Treasury — *Treasury Disrupts Corruption Network Stealing From Venezuela's Food Distribution Program (CLAP)* (25-7-2019; Alex Saab network; Group Grand Ltd; markups); Voz de América — *¿Qué son los CLAP y qué tienen que ver con Alex Saab?*; InsightCrime — *How Criminals Take Advantage of Venezuela's Food Aid System*; Excélsior — *Despensas CLAP, una red de corrupción* (~12M boxes from Mexico, ~$30/box markup, ~$360 million; milk below labeling — Armando.Info). https://home.treasury.gov/system/files/136/2019-07-25-Venezuela-Alex-Saab-Network-SPANISH-TRANSLATION.pdf · https://www.vozdeamerica.com/a/que-son-los-clap-y-que-tienen-que-ver-con-alex-saab-/6277181.html · https://insightcrime.org/news/criminals-take-advantage-venezuelas-food-aid-system/ · https://www.excelsior.com.mx/internacional/despensas-clap-red-corrupcion-maduro-larga-historia
300. Wikipedia (ES) — *Envío de ayuda humanitaria en Venezuela de 2019* (Tienditas bridge blocked with a tanker and containers; 23-2-2019: 14 trucks/280 tonnes, 2 trucks burned; border closure); CNN en Español — *Venezuela bloquea puente en frontera con Colombia* (6-2-2019); Infobae — *Maduro ordenó retirar los contenedores del puente Tienditas* (Dec. 2022). https://es.wikipedia.org/wiki/Env%C3%ADo_de_ayuda_humanitaria_en_Venezuela_de_2019 · https://cnnespanol.cnn.com/2019/02/06/venezuela-bloquea-puente-en-frontera-con-colombia ·

https://www.infobae.com/america/venezuela/2022/12/14/tras-impedir-el-paso-de-la-ayuda-humanitaria-en-2019-maduro-ordeno-retirar-los-contenedores-del-puente-tienditas/

301. NBC News — *Red Cross: $6 million for Ebola fight stolen through fraud* (Nov. 2017; ≥$6 million out of >$124 million; Sierra Leone ~$2.13M collusion with a bank over the exchange rate; Guinea ~$1.17M fraudulent invoicing by a customs-clearance provider; Liberia ~$2.7M inflated prices for relief items, payrolls and incentives); IFRC/ReliefWeb — *IFRC statement on fraud in Ebola operations*; Newsweek — *Red Cross Loses $6 Million In Fraud During Ebola Crisis*. https://www.nbcnews.com/storyline/ebola-virus-outbreak/red-cross-6-million-ebola-fight-stolen-through-fraud-n817701 · https://reliefweb.int/report/sierra-leone/ifrc-statement-fraud-ebola-operations · https://www.newsweek.com/red-cross-outraged-6-million-fraud-loss-during-ebola-outbreak-702231
302. UN News — *UN food agency welcomes review of its food distribution programme in Somalia* (2010; Monitoring Group: "up to half" diverted; contractors/militias/local staff; Adaani family); VOA — *WFP Report says No Evidence of Food Diversion in Somalia* (WFP's denial — Josette Sheeran; own audit found no evidence); ODI — *Somalia food aid diversion*. https://news.un.org/en/story/2010/03/332792-un-food-agency-welcomes-review-its-food-distribution-programme-somalia · https://www.voanews.com/a/un-official-slams-report-on-somalia-food-diversion-89115947/153849.html · https://odi.org/en/insights/somalia-food-aid-diversion/
303. IPS Noticias — *Corrupción-España: investigación de ONG refuerza transparencia* (2007; Intervida: >€45 million diverted according to the Anti-Corruption Prosecutor's Office, ~10,000 donors canceled; Anesvad: president arrested on 5-4-2007 for misappropriation, reported by the workers). https://ipsnoticias.net/2007/04/corrupcion-espana-investigacion-de-ong-refuerza-transparencia/
304. Natural Hazards Center (Univ. Colorado) — *Avoiding the "Second Disaster" of Unwanted Donations* (material convergence; unsolicited donations clog runways and warehouses); TED Ideas — *After a disaster, don't send toys or clothing — send money* (José Holguín-Veras: 50-70% of the goods that arrive are neither needed nor appropriate; the toy-plane case). https://hazards.colorado.edu/news/research-counts/avoiding-the-second-disaster-of-unwanted-donations · https://ideas.ted.com/after-a-disaster-dont-send-toys-or-clothing-send-money-heres-why/
305. WFP — *UN Humanitarian Response Depot (UNHRD)* (depots in Brindisi, Accra, Kuala Lumpur, Panama City and Dubai; dispatch within 48 h; stocks for multiple organizations); UNHRD — *Our Services*; World Economic Forum — *Logistics Emergency Team: 20 years* (Logistics Cluster led by the WFP). https://www.wfp.org/unhrd · https://unhrd.org/our-services · https://www.weforum.org/stories/global-cooperation/logistics-emergency-team-humanitarian-aid-20-years/
306. UNICEF Office of Innovation — *Why blockchain is the future of humanitarian response* (StaTwig: traceability of vaccine vials from depot to health center); WFP Innovation Accelerator — *3 ways blockchain innovation is enhancing humanitarian work* (food traceability; Building Blocks); LSE — *How Blockchain is Contributing to the Humanitarian Sector as of 2025* (IFRC/AID:Tech pilot with the Irish and Lebanese Red Cross). https://www.unicef.org/innovation/stories/why-blockchain-future-humanitarian-response · https://wfpinnovation.medium.com/3-ways-that-blockchain-innovation-is-enhancing-humanitarian-work-e40dd3e85dee · https://blogs.lse.ac.uk/internationaldevelopment/2025/07/10/how-blockchain-is-contributing-to-the-humanitarian-sector-as-of-2025/
307. t2p — product manual (mode wheel; document scanner up to 50 pages in one PDF; SHA-512 fingerprint computed on the device; background certification that resumes when coverage returns; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; device integrity check and hardware attestation; date/time overlay on video; signature with video selfie, liveness challenges and biometrics; .asice ETSI EN 319 162 evidence package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; cash delivery and deduplication, in chapter 49; the donor's euro, in chapter 73.
308. Wikipedia — *2023 Pylos migrant boat disaster* (the Adriana; departure from Libya; ≈750 people aboard; more than 600 dead; 104 survivors; 82 bodies; the coast guard's version); Human Rights Watch — *Greece: Disparities in Accounts of Pylos Shipwreck Underscore the Need for Human Rights Compliant Inquiry* (the authorities' failures; allegations about the tow). https://en.wikipedia.org/wiki/2023_Pylos_migrant_boat_disaster ·

https://www.hrw.org/news/2023/08/03/greece-disparities-accounts-pylos-shipwreck-underscore-need-human-rights-compliant
309. Amnesty International — *Greece: Trial of Pylos shipwreck survivors must not be a shield in the investigation of the alleged responsibilities of the Greek authorities* (May 2024; nine survivors prosecuted; criticism from human rights organizations). https://www.amnesty.org/en/latest/news/2024/05/greece-trial-of-pylos-shipwreck-survivors-must-not-be-a-shield-in-the-investigation-of-the-alleged-responsibilities-of-the-greek-authorities/
310. Refugee Support Aegean — *Pylos Shipwreck: Criminal prosecution for felonies against 17 members of the Coast Guard* (May 26, 2025; captain of the ΠΠΛΣ-920 accused of causing a shipwreck resulting in the death of at least 82 people; complaint on behalf of 40 survivors). https://rsaegean.org/en/pylos-shipwreck-criminal-prosecution-for-felonies-to-17-members-of-the-coast-guard/
311. Global Investigative Journalism Network — *How They Did It: Investigating the Pylos Shipwreck* (journalistic reconstruction with maritime traffic data, survivors' testimonies and material from their phones). https://gijn.org/stories/investigating-pylos-shipwreck/
312. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; only fingerprints on the chain; optional publication; recovery of copies verifiable against the fingerprint). See chapters 13 and 15; human rights evidence, in chapter 20; the position of ships, in chapter 103.
313. Wikipedia — *Border Violence Monitoring Network* (coalition of more than 12 organizations, founded in 2017); Statewatch — *The 'Black Book of Pushbacks': testimonies of pushbacks affecting over 12,000 people* (December 2020; 12,654 people); Balkan Insight — *Study Reports 25,000 Violent Pushbacks at EU Borders* (expanded 2022 edition: more than 3,000 pages, 4 volumes). https://en.wikipedia.org/wiki/Border_Violence_Monitoring_Network · https://statewatch.org/news/2021/january/eu-the-black-book-of-pushbacks-testimonies-of-pushbacks-affecting-over-12-000-people/ · https://balkaninsight.com/2022/12/08/study-reports-25000-violent-pushbacks-at-eu-borders/
314. Iberley — *El TJUE refuerza la tutela judicial efectiva en casos de devoluciones sumarias de migrantes* (case C-136/24 P, judgment of December 18, 2025; Syrian national stripped of his phone and returned to Turkey with Frontex present; adaptation of the burden of proof; Frontex as holder of the documentation). https://www.iberley.es/noticias/el-tjue-refuerza-tutela-judicial-efectiva-casos-devoluciones-sumarias-migrantes-35841
315. Hertfordshire Constabulary — *Three jailed for £4m inheritance fraud* (the Leigh Voysey case and Maureen Renny's will; ~£4.2 million). https://www.herts.police.uk/news/hertfordshire/news/2024/december-2024/three-jailed-for-4m-inheritance-fraud/
316. Cerulli Associates — *Cerulli anticipates $84 trillion in wealth transfers through 2045* (2022; the "great wealth transfer"). https://www.cerulli.com/press-releases/cerulli-anticipates-84-trillion-in-wealth-transfers-through-2045
317. Caring.com / YouGov — *Wills and Estate Planning Study 2025* (only 24% of adults have a will). https://www.caring.com/resources/wills-survey
318. Dutton Gregory / Probate Registry — *Inheritance dispute statistics* (11,362 caveats in 2024, +56% in five years). https://www.duttongregory.co.uk/site/blog/personalnews/inheritance-dispute-statistics
319. Justia — *Will contests: testamentary capacity, undue influence, forgery* (grounds for invalidity). https://www.justia.com/estate-planning/wills/will-contests/
320. CBS News — *QuadrigaCX CEO took password to his grave* (Gerald Cotten; ~$137 million locked). https://www.cbsnews.com/news/gerald-cotten-quadrigacx-ceo-took-password-to-his-grave-leaving-company-unable-to-access-137-million/
321. Chainalysis (via Decrypt) — *~3.7 million bitcoins lost forever* (2020; ~20% of supply). https://decrypt.co/37171/lost-bitcoin-3-7-million-bitcoin-are-probably-gone-forever
322. American Bar Association — *Digital property and estate planning* (RUFADAA; cryptocurrency is lost without the keys). https://www.americanbar.org/groups/real_property_trust_estate/resources/estate-planning/digital-property/
323. European Commission (Interoperable Europe) / e-Estonia — *Estonia's KSI blockchain: the Succession Registry* (first state service on blockchain, 2012). https://e-estonia.com/blockchain-security-control-for-government-registers/
324. eSignGlobal — *Electronic wills legislation and remote notarization* (eIDAS QES; wills often excluded from electronic notarization). https://www.esignglobal.com/blog/electronic-wills-

legislation-remote-notary

325. INE / Gracia Carrión (analysis) — *Custodia compartida supera a la materna en España (INE 2024)* (49.7% of divorces with children, +1.3 points; first time above maternal custody; 12.6% in 2010); APFS — *Separaciones y divorcios en España 2024* (82,991 divorces, +8.2%); INE — *Estadística de nulidades, separaciones y divorcios*. https://graciacarrion.com/es/custodia-compartida-supera-materna/ · https://apfs.es/separaciones-y-divorcios-en-espana-2024-tendencias-y-claves-segun-el-ine/ · https://www.ine.es/dyngs/INEbase/es/operacion.htm?c=estadistica_C&cid=1254736176798&menu=ultiDatos&idp=1254735573206
326. OurFamilyWizard — *Tools for joint custody* (the reference co-parenting application for courts and professionals in all 50 U.S. states, Canada, the United Kingdom, Australia and New Zealand; timestamped messages with no deletion; read confirmations; calendars with history; reports for trials; ToneMeter); Crianza y Crecimiento — *Aplicaciones de coparentalidad aprobadas por tribunales*. https://es.ourfamilywizard.com/homepage · https://crianzaycrecimiento.com/aplicaciones-coparentalidad-aprobadas-tribunales/
327. El Español — *Una abogada, sobre solicitar la custodia compartida: «Después de 10 años, necesitas un plan de parentalidad»* (the parenting plan as the common framework of upbringing beyond the divorce settlement). https://www.elespanol.com/mujer/relaciones/20251214/abogada-solicitar-custodia-compartida-despues-anos-necesitas-plan-parentalidad/1003744039064_0.html
328. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; bilateral record; content on the devices, only fingerprints on the chain; disclosure to the court). See chapters 13 and 15; the legal recording of one's own conversations, in chapter 75.
329. Asian Development Bank (ADB) — *Global trade finance gap expands to $2.5 trillion in 2022* (Trade Finance Gaps, Growth and Jobs Survey). https://www.adb.org/news/global-trade-finance-gap-expands-25-trillion-2022
330. McKinsey — *The multi-billion-dollar paper jam: unlocking trade by digitalizing documentation* (2022; $6.5 billion in savings + new trade). https://www.mckinsey.com/industries/logistics/our-insights/the-multi-billion-dollar-paper-jam-unlocking-trade-by-digitalizing-documentation
331. DCSA / FIT Alliance — *eBL adoption* (real penetration of the electronic bill of lading: 5.7%, January 2025). https://dcsa.org/standards/booking/adoption-guide-booking-and-bill-of-lading
332. Caixin Global — *Dezheng founder jailed over Qingdao port forgery scandal* (Chen Jihong, 23 years; ~12,300M yuan with false receipts). https://www.caixinglobal.com/2018-12-12/dezheng-founder-gets-32-years-in-prison-over-port-forgery-scandal-101358154.html
333. Singapore Police Force / SCMP — *Hin Leong founder OK Lim jailed* (17.5 years, reduced to 13.5 on appeal in March 2026; ~$112 million from HSBC). https://www.police.gov.sg/media-hub/news/2024/20241118_founder_of_hin_leong_trading_sentenced_and_jailed_for_cheating_and_instigating_fc
334. FINMA — *Greensill / Credit Suisse* (2021 insolvency; ~$10 billion in four funds; fictitious invoices). https://www.finma.ch/en/news/2023/02/20230228-mm-greensill/
335. Maersk — *Maersk and IBM to discontinue TradeLens* (consortium shutdown, 2022-2023). https://www.maersk.com/news/articles/2022/11/29/maersk-and-ibm-to-discontinue-tradelens
336. UNCITRAL — *Model Law on Electronic Transferable Records (MLETR): status* (13 jurisdictions; including the United Kingdom). https://uncitral.un.org/en/texts/ecommerce/modellaw/electronic_transferable_records/status
337. Clifford Chance — *UK Electronic Trade Documents Act 2023* (legal equivalence of the electronic bill of lading). https://www.cliffordchance.com/briefings/2023/08/uk-electronic-trade-documents-act-2023--a-further-step-towards-p.html
338. International Chamber of Commerce (ICC) — *Survey shows steady rise of global adoption of electronic bills of lading* (~4 billion documents/day; UCP 600 rules). https://iccwbo.org/news-publications/news/survey-shows-steady-rise-of-global-adoption-of-electronic-bills-of-lading/
339. Financial Times (via S&P Global) — *Wirecard auditor did not independently verify its bank balance* (EY; nonexistent escrow accounts; fabricated statements). https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/wirecard-auditor-did-not-independently-verify-its-bank-balance-8211-financial-times-59219155
340. CNN Business — *Germany bans EY from taking on new audit clients over Wirecard* (APAS: "grave negligence"; €500,000 fine and 2-year ban). https://www.cnn.com/2023/04/03/business/wirecard-ey-ban-germany/index.html
341. Dan McCrum (Financial Times) — *Money Men: A Hot Startup, a Billion-Dollar Fraud, a Fight for the Truth* (2022). https://danmccrum.com/

342. Wikipedia — *Enron scandal* (>$74 billion lost; Arthur Andersen, ~85,000 employees; origin of Sarbanes-Oxley). https://en.wikipedia.org/wiki/Enron_scandal
343. U.S. Securities and Exchange Commission — *Luckin Coffee agrees to pay $180 million penalty to settle accounting fraud charges* (~$311 million in fabricated sales). https://www.sec.gov/newsroom/press-releases/2020-319
344. PCAOB — *AS 2310, The Auditor's Use of Confirmation* (adopted September 2023; direct external confirmation). https://pcaobus.org/oversight/standards/auditing-standards/details/AS2310
345. Association of Certified Fraud Examiners (ACFE) — *Report to the Nations 2024* (5% of revenue; financial statement fraud, a median of $766,000 per case). https://www.acfe.com/-/media/files/acfe/pdfs/rttn/2024/2024-report-to-the-nations.pdf
346. European Commission — *Corporate sustainability reporting (CSRD)* (mandatory limited assurance; EFRAG's ESRS standards). https://finance.ec.europa.eu/financial-markets/company-reporting-and-auditing/company-reporting/corporate-sustainability-reporting_en
347. Accountancy Europe — *Omnibus explained: key changes to the CSRD and CSDDD* (threshold of >1,000 employees and >€450 million; limited assurance remains). https://accountancyeurope.eu/publications/omnibus-explained-key-changes-to-the-csrd-and-csddd/
348. IAASB — *ISSA 5000: General Requirements for Sustainability Assurance Engagements* (2024; first global sustainability assurance standard). https://www.iaasb.org/publications/international-standard-sustainability-assurance-5000-general-requirements-sustainability-assurance
349. Juniper Research (via PR Newswire) — *22% of online ad spend is wasted due to ad fraud in 2023* ($84 billion in 2023 → $172 billion in 2028; mobile 30%). https://www.prnewswire.com/news-releases/new-ad-fraud-study-22-of-online-ad-spend-is-wasted-due-to-ad-fraud-in-2023-according-to-juniper-research-301938050.html
350. Association of National Advertisers (ANA) — *Programmatic Media Supply Chain Transparency Study* (December 2023; 36 cents per dollar; ~$22 billion recoverable; MFA 15%). https://www.ana.net/miccontent/show/id/rr-2023-12-ana-programmatic-media-supply-chain-transparency-study
351. eMarketer — *5 key stats on ad waste, fraud, MFAs and brand risks* (MFA from 15% to 4%). https://www.emarketer.com/content/5-key-stats-ad-waste-fraud-mfas-brand-risks-beyond
352. Imperva — *2024 Bad Bot Report* (bots, 49.6% of internet traffic in 2023; malicious bots, 32%). https://www.imperva.com/company/press_releases/bots-make-up-half-of-all-internet-traffic-globally/
353. Department of Justice (EDNY) — *Two international cybercriminal rings dismantled (Methbot and 3ve)* (November 2018; 1.7 million computers; >$29 million). https://www.justice.gov/usao-edny/pr/two-international-cybercriminal-rings-dismantled-and-eight-defendants-indicted-causing
354. Krebs on Security / White Ops — *Report: $3-5M in ad fraud daily from Methbot* (2016; 300 million views/day). https://krebsonsecurity.com/2016/12/report-3-5m-in-ad-fraud-daily-from-methbot/
355. The Record (Recorded Future) — *"King of Fraud" sentenced to 10 years for Methbot/3ve* (Aleksandr Zhukov, 10-Nov-2021). https://therecord.media/king-of-fraud-sentenced-to-10-years-in-prison-for-role-in-methbot-3ve-botnet
356. IAB Tech Lab — *ads.txt (Authorized Digital Sellers)* (2017; against domain spoofing). https://iabtechlab.com/ads-txt/
357. CBS News / Cheq (Roberto Cavazos, Univ. of Baltimore) — *Influencer marketing fraud costs companies $1.3 billion* (2019). https://www.cbsnews.com/news/influencer-marketing-fraud-costs-companies-1-3-billion/
358. Augustine Fou (FouAnalytics) — independent research on digital advertising fraud. https://www.linkedin.com/pulse/how-do-we-un-fu-digital-advertising-dr-augustine-fou
359. Mandiant / Google Cloud — *M-Trends 2026* (average dwell time of 14 days in 2025; access handoff in 22 seconds). https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026
360. MITRE ATT&CK — *T1070 Indicator Removal / T1070.001 Clear Windows Event Logs* (the purge leaves Event ID 1102). https://attack.mitre.org/techniques/T1070/001/
361. NIST — *SP 800-86: Guide to Integrating Forensic Techniques into Incident Response* (2006; forensic phases and chain of custody). https://csrc.nist.gov/pubs/sp/800/86/final
362. IBM — *Cost of a Data Breach 2025* (global average cost of $4.44 million / $10.22 million in the U.S.; 241 days to identify and contain). https://www.ibm.com/think/x-force/2025-cost-of-a-data-breach-navigating-ai
363. European Commission — *NIS2 Directive (EU 2022/2555)* (notification within 24 h / 72 h; 23 states in infringement proceedings). https://digital-strategy.ec.europa.eu/en/policies/nis2-directive

364. EUR-Lex — *DORA, Regulation (EU) 2022/2554* (applicable since 17-Jan-2025; reporting of serious ICT incidents). https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng
365. U.S. Securities and Exchange Commission — *Cybersecurity disclosure rule* (26-Jul-2023; Form 8-K, item 1.05, within 4 business days). https://www.sec.gov/newsroom/press-releases/2023-139
366. GDPR-info — *Article 33 GDPR* (data breach notification within 72 hours). https://gdpr-info.eu/art-33-gdpr/
367. Global News — *Hamilton cyberattack cost* (insurer denies coverage over incomplete MFA; 18.3M CAD). https://globalnews.ca/news/11313018/hamilton-cyberattack-cost/
368. Richard Bejtlich (TaoSecurity) — *Mandiant global median dwell time* (two consecutive years of deterioration). https://taosecurity.blogspot.com/2026/03/mandiant-global-median-dwell-time.html
369. Carissa Véliz (ed.) — *The Oxford Handbook of Digital Ethics* (Oxford UP, 2022).
370. FedScoop (2024-2025). https://fedscoop.com/ice-wants-more-blockchain-analytics-tech/
371. CJEU — *Judgment of May 4, 2023, case C-300/21 (UI v. Österreichische Post)* (Article 82 GDPR: three cumulative conditions — infringement, harm and causal link; the mere infringement confers no right to compensation; non-material harm is compensable with no minimum threshold of seriousness; quantification under national law subject to equivalence and effectiveness); analyses by Cuatrecasas, finReg360 and Confilegal. https://www.cuatrecasas.com/es/spain/art/indemnizacion-de-danos-morales-por-infraccion-del-rgpd · https://finreg360.com/alerta/el-tjue-dicta-sentencia-sobre-el-derecho-de-indemnizacion-por-tratamiento-de-datos-personales/ · https://confilegal.com/20230913-el-tjue-admite-que-los-danos-no-materiales-por-el-uso-ilicito-de-datos-personales-merecen-una-compensacion-economica/ · https://www.iberley.es/noticias/tjue-mero-incumplimiento-normativa-proteccion-datos-da-lugar-responsabilidad-pero-solo-dano-probado-resulta-indemnizable-33163
372. GDPR (Arts. 33-34: breach notification to the authority within 72 h and to those affected in case of high risk; Art. 82: right to compensation); Have I Been Pwned as a public leak-checking tool. See chapter 57 for the company's side.
373. t2p — product manual (certified screenshot and web page capture — full image + source code; audio recording of one's own conversations; certify any file; projects for an incident diary; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; SIM swapping, in chapter 68; the scam with leaked data, in chapter 85.
374. Wikipedia — *2024 Delta Air Lines disruption* / TechTarget — *CrowdStrike outage explained* (19-7-2024: faulty Falcon Sensor update; ~8.5 million Windows machines; the biggest IT outage in history; impact on airlines, hospitals, banks, emergency services); CNBC — *Delta, CrowdStrike sue each other over widespread IT outage* (Delta: >7,000 flights canceled over five days, >1.3 million passengers; October 2024 suit for ~$550 million; cross-suits). https://en.wikipedia.org/wiki/2024_Delta_Air_Lines_disruption · https://www.techtarget.com/whatis/feature/Explaining-the-largest-IT-outage-in-history-and-whats-next · https://www.cnbc.com/2024/10/25/delta-suit-against-crowdstrike-after-it-outage-caused-cancellations.html
375. Regulation (EU) 2022/2554 (DORA — digital operational resilience of the financial sector: third-party ICT risk management, incident reporting; applicable since 11-1-2025) and the NIS2 Directive (expanded incident notification obligations). See chapter 57 for the security context and notification deadlines.
376. t2p — product manual (certify any file — only the fingerprint travels; web page capture with full image + source code; certified screenshot and screen recording; projects to group an incident's file; background certification; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the records of an attack, in chapter 57; the insurance claim, in chapter 28.
377. CNBC — *This $80 million art scam fooled experts* (Knoedler gallery; Pei-Shen Qian; <$9,000 per work). https://www.cnbc.com/2018/08/16/think-you-can-spot-a-fraud-this-80-million-art-scam-fooled-experts.html
378. The Art Newspaper — *How the detection of titanium white gave away Beltracchi* (2012; the impossible pigment). https://www.theartnewspaper.com/2012/12/01/how-the-detection-of-titanium-white-gave-away-beltracchi

379. OCCRP — *India sentences notorious antiquities trafficker from Manhattan* (Subhash Kapoor; >2,500 pieces, ~$143 million; the Chola Shiva). https://www.occrp.org/en/news/india-sentences-notorious-antiquities-trafficker-from-manhattan
380. Manhattan District Attorney's Office — *Antiquities Trafficking Unit* (~6,000 antiquities from 29 countries, >$460 million recovered). https://manhattanda.org/d-a-bragg-announces-return-of-antiquities-to-spain-and-cyprus/
381. The Conversation — *The Gurlitt case: how a routine customs check uncovered a Nazi-era art hoard* (1,406 works). https://theconversation.com/the-gurlitt-case-how-a-routine-customs-check-uncovered-a-sensational-nazi-era-art-hoard-105604
382. U.S. Department of State — *Washington Conference Principles on Nazi-Confiscated Art* (1998; 44 countries). https://www.state.gov/washington-conference-principles-on-nazi-confiscated-art
383. UNESCO — *1970 Convention against the illicit trafficking of cultural property* (~149 states). https://www.unesco.org/en/fight-illicit-trafficking/about
384. INTERPOL — *Cultural heritage crime: the issues* (illicit trade of $225 million to $6 billion/year). https://www.interpol.int/en/Crimes/Cultural-heritage-crime/The-issues-cultural-property
385. Art Loss Register — *~700,000 registered objects; founded by Julian Radcliffe (1991); >$100 million recovered*. https://www.artloss.com/
386. CNN — *The forensic archaeologist hunting trafficked antiquities* (Christos Tsirogiannis; "the papers lie"). https://edition.cnn.com/2025/12/13/style/forensic-archaeologist-trafficked-antiquities-scli-intl
387. UNESCO — *New UNESCO report warns that Generative AI threatens Holocaust memory* (June 2024, with the World Jewish Congress; ChatGPT and the false "Holocaust by drowning"; testimony quotes fabricated by Bard; oversimplification; recommendations to technology companies and education systems); press release in Spanish (PDF). https://www.unesco.org/en/articles/new-unesco-report-warns-generative-ai-threatens-holocaust-memory · https://articles.unesco.org/sites/default/files/medias/fichiers/2024/06/PR_New_UNESCO_report_warns_that_Ge
388. Maldita.es — *El papa, Margarita Salas y escenas de la Guerra Civil generadas con IA, en portada: cómo el uso de imágenes falsas en webs y revistas puede desinformar* (February 25, 2025; synthetic images of real events and people published without a label; alteration of historical perception and distrust in the media). https://maldita.es/malditatecnologia/20250225/imagenes-IA-acontecimientos-reales-desinformar/
389. Forbes — *Internet History Hacked, Wayback Machine Down — 31 Million Passwords Stolen* (October 2024; 7 TB leaked; GitLab token exposed since late 2022; simultaneous DDoS on the Wayback Machine); NPR — *Hackers steal information from 31 million Internet Archive users*; Hispasec — *Internet Archive sufre su segunda brecha* (more than 800,000 support tickets exposed). https://www.forbes.com/sites/daveywinder/2024/10/10/internet-hacked-wayback-machine-down-31-million-passwords-stolen/ · https://www.npr.org/2024/10/20/nx-s1-5159000/internet-archive-hack-leak-wayback-machine · https://unaaldia.hispasec.com/2024/10/internet-archive-sufre-su-segunda-brecha-de-seguridad-en-octubre-de-2024.html
390. Infobae — *La Ley de Memoria Democrática impulsa la exhumación de fosas y la reparación de 9.000 víctimas: todavía quedan más de 10.000 cuerpos por recuperar* (November 18, 2025; 8,941 remains located and recovered; only 70 identified and returned; Integrated Map of Graves: 5,848 graves, ≈1,500 opened; State DNA Bank); El Español — *La nueva Ley de Memoria entra en vigor: el Gobierno quiere exhumar otros 25.000 cadáveres* (October 2022). https://www.infobae.com/espana/2025/11/18/la-ley-de-memoria-democratica-impulsa-la-exhumacion-de-fosas-y-la-reparacion-de-9000-victimas-todavia-quedan-mas-de-10000-cuerpos-por-recuperar/ · https://www.elespanol.com/espana/20221020/nueva-ley-memoria-gobierno-quiere-exhumar-cadaveres/712178865_0.html
391. t2p — product manual (SHA-512 fingerprint at the source; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; batch manifests; content in each archive's systems, only fingerprints on the chain). See chapters 13 and 15; official statements, in chapter 86; war crimes evidence, in chapter 20.
392. EUR-Lex — *Directive (EU) 2019/1937 on the protection of whistleblowers* (confidentiality of identity; secure channels; deadlines). https://eur-lex.europa.eu/eli/dir/2019/1937/oj/eng
393. eucrim — *ECJ orders several member states to financial penalties for failing to transpose the Whistleblowers Directive* (March 2025; ~€39 million; Germany €34 million; Estonia €500,000 + €1,500/day). https://eucrim.eu/news/ecj-ordered-several-member-states-to-financial-penalties-for-failing-to-transpose-whistleblowers-directive/

394. Freedom of the Press Foundation — *SecureDrop* (more than 60 newsrooms; Tor hidden service; architectural source protection). https://securedrop.org/overview/
395. ICIJ — *Panama Papers* (11.5 million documents; 2.6 TB; >350 journalists; the source "John Doe"). https://www.icij.org/investigations/panama-papers/
396. OCCRP — *John Doe: "The Revolution Will Be Digitized"* (the Panama Papers source's manifesto). https://www.occrp.org/en/investigation/panama-papers-source-john-doe-the-revolution-will-be-digitized
397. Errata Security (Robert Graham) — *How The Intercept outed Reality Winner* (printer tracking dots; serial 29535218; 9-May-2017, 6:20). https://blog.erratasec.com/2017/06/how-intercept-outed-reality-winner.html
398. Freedom of the Press Foundation — *Reality Winner sentenced to 63 months* (the longest U.S. sentence for a press source). https://freedom.press/issues/whistleblower-reality-winner-sentenced-longest-prison-sentence-history-federal-leak-cases/
399. Electronic Frontier Foundation (EFF) — *List of printers which do or do not display tracking dots* (machine identification code). https://www.eff.org/pages/list-printers-which-do-or-do-not-display-tracking-dots
400. Wikipedia — *Stylometry / Adversarial stylometry* (authorship attribution with 80-95% accuracy). https://en.wikipedia.org/wiki/Stylometry
401. Wikipedia — *Frances Haugen* (the Facebook leak; chose to reveal her identity). https://en.wikipedia.org/wiki/Frances_Haugen
402. Carissa Véliz — *The Ethics of Privacy and Surveillance* (Oxford UP, 2024).
403. Nature (Richard Van Noorden) — *More than 10,000 research papers were retracted in 2023 — a new record* (Hindawi/Wiley, >8,000). https://www.nature.com/articles/d41586-023-03974-8
404. Elisabeth Bik, Arturo Casadevall and Ferric Fang — *The Prevalence of Inappropriate Image Duplication in Biomedical Research Publications* (mBio, 2016; 20,621 papers, 3.8%). https://journals.asm.org/doi/10.1128/mbio.00809-16
405. American Society for Microbiology / mBio — *ImageTwin pilot* (2,627 manuscripts; 410 with image problems; 6 acceptances revoked). https://journals.asm.org/doi/10.1128/mbio.01990-25
406. U.S. Department of Justice — *Dana-Farber Cancer Institute agrees to pay $15M to settle fraud allegations* (images/data misrepresented in NIH grants). https://www.justice.gov/opa/pr/dana-farber-cancer-institute-agrees-pay-15m-settle-fraud-allegations-related-scientific
407. Stanford Board of Trustees — *Report on the scientific work of Marc Tessier-Lavigne* (2023; the president's resignation). https://boardoftrustees.stanford.edu/news/board-trustees-statement-release-report-and-announcements
408. American Physical Society — *The Schön scandal* (Beasley committee; misconduct on 16 of 24 charges). https://www.aps.org/apsnews/2022/08/september-2002-schon-scandal-report
409. Retraction Watch (Ivan Oransky and Adam Marcus) — *The new retraction record holder* (Joachim Boldt, 184 retractions, July 2023). https://retractionwatch.com/2023/07/12/the-new-retraction-record-holder-is-a-german-anesthesiologist-with-184/
410. Ben Goldacre et al. (COMPare, University of Oxford) — *COMPare: outcome-switching in clinical trials* (Trials, 2019; 76.3% of primary outcomes properly reported). https://pmc.ncbi.nlm.nih.gov/articles/PMC6375128/
411. Center for Open Science (Brian Nosek) — *Registered Reports* (from 91 to more than 300 journals; the protocol before the results). https://www.cos.io/initiatives/registered-reports
412. eLife — *Reproducibility Project: Cancer Biology* (Amgen, 11% of 53 studies; replicated effects 85% smaller). https://elifesciences.org/articles/71601
413. Wikipedia (with BBC / AFP / Japan Times) — *Sogen Kato* (mummified 32 years; survivor's pension collected; ~9.5 million yen). https://en.wikipedia.org/wiki/Sogen_Kato
414. France 24 — *Over 230,000 listed centenarians missing in Japan* (234,354 unconfirmed; 77,118 aged 120+). https://www.france24.com/en/20100910-over-230000-listed-centenarians-missing-presumed-dead-government-report-japan
415. SSA Office of the Inspector General — *State death report discrepancies led to $327M in improper payments* (2025) and *6.5 million holders aged ≥112 with no date of death in the Numident* (2015). https://oig.ssa.gov/news-releases/2025-03-19-state-death-report-discrepancies-led-to-327m-in-improper-payments-and-cost-ssa-millions-of-dollars-in-unnecessary-workloads/
416. Reuters (via HuffPost) — *Greece stops pensions and benefits for 200,000 people* (~9,000 "centenarians"; savings of up to €800 million/year). https://www.huffpost.com/entry/greece-stops-pensions-benefits-200000-people-fraudulently_n_1451423

417. Moneyweb — *SASSA explains why it paid R140m in grants to 75,000 dead people* (2025; 75% of the fraud by officials, according to the IEJ). https://www.moneyweb.co.za/news/south-africa/sassa-explains-why-it-paid-r140m-in-grants-to-75-000-dead-people/
418. Outlook Money — *Jeevan Pramaan: digital life certificate* (~19 million certificates; ~11.5M via facial authentication). https://www.outlookmoney.com/retirement/digital-life-certificate-over-14-million-super-seniors-generated-online-jeevan-pramaan-patra-in-fy25-26
419. The Wire (India) — *Digital Exclusion: Poor, Elderly Face the Brunt of Aadhaar-Based Authentication Errors* (worn fingerprints; pensions denied; fingerprint made non-mandatory). https://m.thewire.in/article/rights/digital-exclusion-poor-elderly-face-the-brunt-of-aadhaar-based-authentication-errors
420. GOV.UK (DWP) — *State Pension underpayments (operation LEAP)* (119,050 underpayments, £736 million; >15,000 already dead). https://www.gov.uk/government/publications/state-pension-underpayments-progress-on-cases-reviewed-to-30-september-2024/state-pension-underpayments-progress-on-cases-reviewed-to-30-september-2024
421. International Social Security Association (ISSA) — *Proof of life in international social security* and the VIVESS good practice (Spanish pensioners abroad). https://www.issa.int/analysis/proof-life-international-social-security-progress-and-persistent-challenges
422. Consultor Jurídico (Conjur, Brazil) — *Deepfakes, provas digitais e fraudes no INSS: o colapso da confiança* (2025). https://www.conjur.com.br/2025-mai-28/deepfakes-provas-digitais-e-fraudes-no-inss-colapso-da-confianca/
423. CNN en Español — *Una mujer cree que estafadores clonaron la voz de su hija en un falso secuestro* (the Jennifer DeStefano case, Arizona; call of January 20 around 16:55; 15-year-old daughter at a competition; $1 million ransom lowered to $50,000; "the voice was so real"; later testimony before a U.S. Senate subcommittee); La Opinión — *Utilizan voz recreada por inteligencia artificial para fingir el secuestro de una niña y pedir rescate*. https://cnnespanol.cnn.com/2023/04/29/una-mujer-cree-que-los-estafadores-clonaron-la-voz-de-su-hija-en-un-falso-secuestro-trax/ · https://laopinion.com/2023/04/29/utilizan-voz-recreada-por-inteligencia-artificial-para-fingir-el-secuestro-de-una-nina-y-pedir-rescate/
424. FBI — *'Virtual Kidnapping' Extortion Calls on the Rise* (the scam's mechanics; international numbers; blocking contact; ransom by wire; recommendations: calm, contact by another route, questions only the real person can answer); FBI Albuquerque — *Be on the Lookout for Virtual Kidnapping Scams*. https://www.fbi.gov/contact-us/field-offices/sanantonio/news/press-releases/virtual-kidnapping-extortion-calls-on-the-rise · https://www.fbi.gov/contact-us/field-offices/albuquerque/news/press-releases/fbi-media-alert-be-on-the-lookout-for-virtual-kidnapping-scams
425. FBI / IC3 — *Criminals Using Altered Proof-of-Life Media to Extort Victims in Virtual Kidnapping for Ransom Scams* (2025 alert: photos and videos altered or AI-generated as false "proof of life," often from public images; inconsistencies such as missing tattoos or scars; report to ic3.gov with the material received); Axios — *AI-generated videos, photos used in new virtual kidnapping scams, FBI warns* (December 5, 2025). https://www.fbi.gov/investigate/cyber/alerts/2025/criminals-using-altered-proof-of-life-media-to-extort-victims-in-virtual-kidnapping-for-ransom-scams · https://www.axios.com/2025/12/05/ai-videos-photos-scam-virtual-fbi-victim-ransom-kidnapping
426. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; proof of life with challenge and freshness — an unpredictable nonce answered and sealed to the minute; voluntary trail with chosen moments; no central registry; revocability). See chapters 13, 15 and 64; the cloned voice, in chapter 68.
427. La Moncloa / Ministry of the Interior — *Las Fuerzas de Seguridad investigaron la desaparición de 16.147 personas en 2024* (+6%; 72% of cases resolved in under a week; CNDES, a pioneer in Europe); Ministry of the Interior — *Las Fuerzas de Seguridad investigaron la desaparición de 16.024 personas en 2025 y esclarecieron el 91 por ciento de las denuncias* (25,086 reports; 928 alerts published, 132 active; 11,461 unique episodes, an average of 3 reports for repeat cases). https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/interior/paginas/2025/170325-marlaska-plan-personas-desaparecidas.aspx · https://www.interior.gob.es/opencms/es/detalle/articulo/Las-Fuerzas-de-Seguridad-investigaron-la-desaparicion-de-16.024-personas-en-2025-y-esclarecieron-el-91-por-ciento-de-las-denuncias/
428. Wikipedia (synthesis of the *McLaren Report*) — *McLaren Report* (≥643 "disappeared" positives; male DNA in female samples). https://en.wikipedia.org/wiki/McLaren_Report

429. World Anti-Doping Agency — *Independent investigation confirms Russian state manipulation of doping control* (McLaren, Independent Person; >1,000 athletes). https://www.wada-ama.org/en/news/wada-statement-independent-investigation-confirms-russian-state-manipulation-doping-control
430. CNN — *WADA McLaren report: tampered samples* (44 samples with marks on the caps, 12 from medalists; Part 2, 2016). https://www.cnn.com/2016/12/09/sport/wada-russia-doping-mclaren-report
431. Chemistry World — *Anti-doping scientists expose cheating Russian athletes* (8 samples with physiologically impossible salt). https://www.chemistryworld.com/news/anti-doping-scientists-expose-cheating-russian-athletes/2500159.article
432. Republik — *Bad bottles* (the "tamper-proof" Berlinger bottles, breached). https://www.republik.ch/2018/02/21/bad-bottles
433. World Anti-Doping Agency — *Athlete Biological Passport* (indirect detection; hematological module 2009, steroidal 2014). https://www.wada-ama.org/en/athlete-biological-passport
434. ESPN — *Pellizotti banned in landmark biological-passport ruling* (2 years without a positive; first case upheld by CAS). https://www.espn.com/olympics/cycling/news/story?id=6192949
435. Europol — *Results from the largest football match-fixing investigation in Europe (Operation VETO)* (>380 matches; 425 people; >15 countries). https://www.europol.europa.eu/media-press/newsroom/news/update-results-largest-football-match-fixing-investigation-in-europe
436. Sportradar — *Betting Corruption and Match-Fixing 2024* (>850,000 matches; 1,108 suspicious; the UFDS system). https://sportradar.com/content-hub/news/sportradar-reports-notable-decline-in-match-fixing-in-2024/
437. The Washington Post — *Grigor Sargsyan, "the Maestro," and tennis match-fixing* (5 years; ~180 tennis players from >35 countries; ~€7 million). https://www.washingtonpost.com/world/interactive/2023/tennis-match-fixing-itf-grigor-sargsyan/
438. U.S. Department of Justice / SDNY (via Music Business Worldwide) — *North Carolina man pleads guilty to AI-assisted music streaming fraud* (Michael Smith; $8.09 million forfeiture). https://www.musicbusinessworldwide.com/streaming-fraud-man-who-pocketed-8m-using-hundreds-of-thousands-of-ai-songs-streamed-billions-of-times-by-bots-pleads-guilty/
439. Rolling Stone — *Inside the AI streaming-fraud case* (from 1,040 to 10,000 bot accounts; 661,440 plays/day; the emails). https://www.rollingstone.com/music/music-features/streaming-fraud-fake-streams-mike-smith-1235500686/
440. Beatdapp / Music Business Worldwide — *Streaming fraud costs the global music industry ~$2bn a year* (≥10% of activity; an interested party's estimate). https://www.musicbusinessworldwide.com/streaming-fraud-costs-the-global-music-industry-2bn-a-year-according-to-beatdapp-now-its-partnering-with-beatport-to-combat-the-trend/
441. Spotify / Music Business Worldwide — *The 1,000-stream royalty threshold* (announced Nov-2023; in force Apr-2024). https://www.musicbusinessworldwide.com/heres-exactly-spotifys-new-1000-annual-streams-royalty-policy-works123/
442. The MLC / Billboard — *Mechanical Licensing Collective receives over $424 million in unmatched royalties* (the "black box"). https://www.billboard.com/pro/mechanical-licensing-collective-receives-over-424-million-unmatched-royalties-streaming-services/
443. IFPI — *Global Music Report 2026* ($31.7 billion; streaming >$22 billion, 69.6%; "theft, plain and simple"). https://www.ifpi.org/global-music-report-2026-global-recorded-music-revenues-grow-6-4-as-record-companies-drive-innovation/
444. IFPI (Victoria Oakley) — *Brazilian authorities disrupt JustAnotherPanel* ("fake streams… erode trust in our system"). https://www.ifpi.org/brazilian-authorities-launch-largest-ever-disruption-operation-against-major-streaming-fraud-platform-justanotherpanel/
445. Federal Trade Commission — *FTC brings first-ever cases under the BOTS Act* (2021; >150,000 tickets; >$31 million in penalties). https://www.ftc.gov/news-events/news/press-releases/2021/01/ftc-brings-first-ever-cases-under-bots-act
446. Federal Trade Commission — *FTC sues Live Nation/Ticketmaster* (2025; 5 brokers, 6,345 accounts, 246,407 tickets; "turn a blind eye as a matter of policy"). https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-sues-live-nation-ticketmaster-engaging-illegal-ticket-resale-tactics-deceiving-artists-consumers
447. Wikipedia — *Taylor Swift–Ticketmaster controversy* (the Eras tour; 3.5 billion requests; ~12 million "entities" with bots). https://en.wikipedia.org/wiki/Taylor_Swift%E2%80%93Ticketmaster_controversy

448. FBI Internet Crime Complaint Center (IC3) — *2023 Internet Crime Report* (SIM swapping: 1,075 complaints, $48.8 million in losses). https://www.ic3.gov/AnnualReport/Reports
449. NPR — *The AI deepfake Biden robocall and the FCC* (New Hampshire; Steve Kramer; $6 million fine). https://www.npr.org/2024/05/23/nx-s1-4977582/fcc-ai-deepfake-robocall-biden-new-hampshire-political-operative
450. Federal Communications Commission — *FCC makes AI-generated voices in robocalls illegal* (Declaratory Ruling, February 2024). https://www.fcc.gov/document/fcc-makes-ai-generated-voices-robocalls-illegal
451. Federal Communications Commission — *FCC settles Lingo Telecom robocall investigation* ($1 million; STIR/SHAKEN + know your customer and your upstream provider). https://www.fcc.gov/document/fcc-settles-lingo-telecom-robocall-investigation
452. TransNexus — *SHAKEN statistics, April 2025* (only 42.5% of calls signed; the non-IP gap). https://transnexus.com/blog/2025/shaken-statistics-april/
453. Cifas — *Fraudscape 2025* (+1,055% unauthorized SIM swaps in the United Kingdom). https://www.cifas.org.uk/newsroom/huge-surge-see-sim-swaps-hit-telco-and-mobile
454. SecurityWeek — *T-Mobile coughed up $33 million in SIM-swap lawsuit* (Josh Jones; ~$38 million in cryptocurrency). https://www.securityweek.com/t-mobile-coughed-up-33-million-in-sim-swap-lawsuit/
455. Federal Trade Commission — *Winners of the Voice Cloning Challenge (OriginStory)* (authenticating the human origin of the voice at the point of creation). https://www.ftc.gov/news-events/news/press-releases/2024/04/ftc-announces-winners-voice-cloning-challenge
456. GSMA — *Open Gateway / Mobile Identity APIs (SIM Swap API)* (the defense migrates to the network layer). https://www.gsma.com/solutions-and-impact/technologies/mobile-identity/mobile-identity-and-data-apis/
457. The Next Web / The Wall Street Journal — *Fraudsters deepfake a CEO's voice* (2019; ~$243,000; the first known synthetic-voice fraud). http://thenextweb.com/news/fraudsters-deepfake-ceos-voice-to-trick-manager-into-transferring-243000
458. Payment Systems Regulator (United Kingdom) — *APP fraud reimbursement protections* (mandatory regime since 7-10-2024, the world's first; individuals, microenterprises and charities; Faster Payments and CHAPS; £85,000 limit — lowered from 415,000 to align with the FSCS; 50/50 cost sharing between sending and receiving provider; reimbursement in 5 business days or up to 35 with investigation; replaces the voluntary CRM code); analyses by Freshfields, A&O Shearman and Thomson Reuters. https://www.psr.org.uk/information-for-consumers/app-fraud-reimbursement-protections/ · https://www.freshfields.com/en/our-thinking/briefings/2024/09/authorised-push-payment-fraud-a-new-mandatory-reimbursement-regime-for-uk-psps/ · https://www.aoshearman.com/en/insights/ao-shearman-on-fintech-and-digital-assets/the-uks-authorised-push-payment-app-fraud-reimbursement-scheme · https://www.thomsonreuters.com/en-us/posts/investigation-fraud-and-risk/app-fraud-uk/
459. PSD2 (Directive (EU) 2015/2366, transposed in Spain by RDL 19/2018): in unauthorized payment transactions, refund by the provider unless there is fraud or gross negligence by the customer; burden of proof of correct authentication on the provider; strong customer authentication (SCA); the European payments reform (PSD3/PSR) in progress.
460. t2p — product manual (audio recording of one's own conversations; certified screenshot and web page capture — full image + source code; certify any file in batches; projects for an incident file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline; only fingerprints on the chain). See chapters 13 and 15; the cloned voice and the SIM, in chapter 68; the impersonated CFO, in chapter 87; the personal black box, in chapter 75.
461. U.S. Senate Committee on Armed Services — *Inquiry into Counterfeit Electronic Parts in the Department of Defense Supply Chain* (2012; >1 million parts; >70% traced to China; THAAD; Hong Dark). https://www.armed-services.senate.gov/press-releases/senate-armed-services-committee-releases-report-on-counterfeit-electronic-parts
462. U.S. Government Accountability Office — *DOD Supply Chain: Suspect Counterfeit Electronic Parts on Internet Platforms* (GAO-12-375; 16 of 16 parts fake). https://www.gao.gov/products/gao-12-375
463. Center for Public Integrity — *Counterfeit chips plague Pentagon weapons systems* (VisionTech; the polish on the chips; McCloskey, 38 months). https://publicintegrity.org/national-security/counterfeit-chips-plague-pentagon-weapons-systems/
464. Air & Space Forces Magazine — *C-27J and C-130J parts linked to counterfeit mills* (L-3; blank screen; 27% defects). https://www.airandspaceforces.com/c-27jandc-

130jpartslinkedtocounterfeitmills/
465. Semiconductor Industry Association — *Winning the battle against counterfeit semiconductor products* (>$7.5 billion/year; ~11,000 jobs). https://www.semiconductors.org/resources/winning-the-battle-against-counterfeit-semiconductor-products/
466. ERAI — *2024 Annual Report* (1,055 suspect parts; the highest since 2015). https://www.erai.com/erai_blog/3187/_2024_annual_report
467. U.S. Department of Justice (NDCA) — *Steve H.S. Kim sentenced for selling $3.5M in counterfeit and substandard parts* (42 months; nuclear submarine, missile system). https://www.justice.gov/usao-ndca/pr/california-resident-steve-hs-kim-sentenced-selling-35m-counterfeit-and-substandard
468. SAE International — *AS6171: test methods for suspect/counterfeit EEE parts* (ISO/IEC 17025 accredited laboratories). https://www.sae.org/standards/content/as6171/7/
469. CALCE (University of Maryland) — *Counterfeit electronics detection and avoidance* (a recycled chip can look identical to the authentic one). https://calce.umd.edu/counterfeit-electronics-detection-and-avoidance
470. Defense Microelectronics Activity (DMEA) — *Trusted Foundry / Trusted IC* (trusted suppliers with a controlled pedigree; 2003, with the NSA). https://www.dmea.osd.mil/TrustedIC.aspx
471. Wikipedia — *XZ Utils backdoor* (CVE-2024-3094; CVSS 10.0; Andres Freund, Microsoft engineer and PostgreSQL developer, investigating ~500 ms of extra SSH latency; report to Openwall on March 28-29, 2024; "Jia Tan's" social engineering campaign with fake accounts pressuring the maintainer; code hidden in build-to-host.m4 and binary test files inside the tarballs — invisible to repository review; potential greater than SolarWinds; caught before mass deployment); Akamai — *XZ Utils Backdoor — Everything You Need to Know*; JFrog — *XZ Backdoor Attack CVE-2024-3094*; Datadog Security Labs — *The XZ Utils backdoor*. https://en.wikipedia.org/wiki/XZ_Utils_backdoor · https://www.akamai.com/blog/security-research/critical-linux-backdoor-xz-utils-discovered-what-to-know · https://jfrog.com/blog/xz-backdoor-attack-cve-2024-3094-all-you-need-to-know/ · https://securitylabs.datadoghq.com/articles/xz-backdoor-cve-2024-3094/
472. Palo Alto Unit 42 — *"Shai-Hulud" Worm Compromises npm Ecosystem in Supply Chain Attack* (September 14, 2025; the first self-replicating worm on npm; >500 packages; @ctrl/tinycolor with >2.2 million weekly downloads; CrowdStrike packages; AWS/GCP/Azure credential harvesting with TruffleHog; persistence via GitHub Actions; propagation with the maintainers' publishing rights); CISA — *Widespread Supply Chain Compromise Impacting npm Ecosystem* (alert of September 23, 2025); StepSecurity — *Shai-Hulud: Self-Replicating Worm Compromises 500+ NPM Packages*; Datadog — *The Shai-Hulud 2.0 npm worm* (second wave, November). https://unit42.paloaltonetworks.com/npm-supply-chain-attack/ · https://www.cisa.gov/news-events/alerts/2025/09/23/widespread-supply-chain-compromise-impacting-npm-ecosystem · https://www.stepsecurity.io/blog/ctrl-tinycolor-and-40-npm-packages-compromised · https://securitylabs.datadoghq.com/articles/shai-hulud-2.0-npm-worm/
473. Sigstore — *Rekor: Software Supply Chain Transparency Log* (public append-only log; >101 million entries from >33,000 projects: Kubernetes, CPython, Istio, urllib3; default signing method for Kubernetes, npm, Maven and PyPI; provenance attestations on npm and PyPI); OpenSSF / Trail of Bits — *Catching malicious package releases using a transparency log* (December 2025; monitoring Rekor to detect fraudulent releases with stolen credentials). https://docs.sigstore.dev/logging/overview/ · https://www.sigstore.dev/ · https://openssf.org/blog/2025/12/19/catching-malicious-package-releases-using-a-transparency-log/ · https://blog.trailofbits.com/2025/12/12/catching-malicious-package-releases-using-a-transparency-log/
474. OJEU / BOE — *Regulation (EU) 2024/2847, of October 23, 2024* (Cyber Resilience Act: horizontal cybersecurity requirements for products with digital elements; in force in December 2024; notification obligations from September 11, 2026; full application on December 11, 2027; security across the whole life cycle); Atalaya GRC — *CRA: fechas y obligaciones*; Red Seguridad — *Cyber Resilience Act: una nueva era para los fabricantes*. https://www.boe.es/buscar/doc.php?id=DOUE-L-2024-81720 · https://atalayagrc.com/normativas/cra/ · https://www.redseguridad.com/especialidades-tic/normativa-y-certificacion/cyber-resilience-act-cra-una-nueva-era-para-los-fabricantes_20260204.html
475. t2p — product manual (SHA-512 fingerprint; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; version fingerprints and batch

manifests; neutral, transversal anchoring, complementary to sector transparency logs). See chapters 13 and 15; the hardware pedigree, in chapter 70; the unerasable records, in chapter 57.

476. Agnew et al. (PLOS ONE) — *Estimating the worldwide extent of illegal fishing* (11-26 million tonnes; $10,000-$23,500 million/year). https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0004570
477. Oceana — *Oceana study reveals seafood fraud nationwide* (33% mislabeled; "red snapper" 87%). https://oceana.org/reports/oceana-study-reveals-seafood-fraud-nationwide/
478. Global Fishing Watch / Nature — *Satellite mapping reveals extensive industrial activity at sea* (~75% of vessels with no public tracking). https://globalfishingwatch.org/press-release/new-research-harnesses-ai-and-satellite-imagery-to-reveal-the-expanding-footprint-of-human-activity-at-sea/
479. European Commission — *Illegal fishing (IUU): catch certificates and CATCH* (Reg. (EC) 1005/2008; CATCH mandatory since 10-Jan-2026). https://oceans-and-fisheries.ec.europa.eu/fisheries/rules/illegal-fishing_en
480. Global Witness — *The ITSCI Laundromat* (up to 80-90% of "clean" mineral from illicit sources). https://www.globalwitness.org/en/campaigns/natural-resource-governance/itsci-laundromat/
481. Kimberley Process / World Diamond Council — *The Kimberley Process fact sheet* (from ~15% in the 90s to <1%; ~99% of production). https://www.worlddiamondcouncil.org/wp-content/uploads/2022/02/The-Kimberley-Process-Fact-Sheet.pdf
482. Environmental Justice Foundation — *Thailand's Seafood Slaves* (slavery on board; the island of Benjina). https://ejfoundation.org/reports/thailands-seafood-slaves
483. Science Advances (Park et al.) — *Illuminating Dark Fishing Fleets in North Korea* (>900 vessels; >164,000 t; >$440 million). https://www.science.org/doi/10.1126/sciadv.abb1197
484. Ledger Insights — *Everledger bankruptcy* (the risk of centralizing diamond traceability in a single company). https://www.ledgerinsights.com/everledger-bankruptcy-esg-blockchain-traceability/
485. The Africa Report — *Rubaya coltan mine at the heart of M23 financing* (~$800,000/month for an armed group). https://www.theafricareport.com/375904/drc-rwanda-rubaya-coltan-mine-at-the-heart-of-m23-financing/
486. ProPublica / NPR (Justin Elliott and Laura Sullivan) — *How the Red Cross raised half a billion dollars for Haiti and built six homes* (2015). https://www.propublica.org/article/how-the-red-cross-raised-half-a-billion-dollars-for-haiti-and-built-6-homes
487. Federal Trade Commission — *FTC, all 50 states and D.C. charge four cancer charities with bilking over $187 million* (the Reynolds family; <3% to patients). https://www.ftc.gov/news-events/news/press-releases/2015/05/ftc-all-50-states-dc-charge-four-cancer-charities-bilking-over-187-million-consumers
488. FBI Internet Crime Complaint Center — *PSA: charity and disaster-relief fraud* (2024; ~$96 million; AI boosts apparent legitimacy). https://www.ic3.gov/PSA/2025/PSA250116
489. New York Attorney General (Letitia James) — *Pennies for Charity 2025* (fundraisers keep 25%; <50% in half of campaigns). https://ag.ny.gov/press-release/2025/attorney-general-james-releases-annual-pennies-charity-report
490. CBS News — *CBS News investigates Wounded Warrior Project spending* (from $1.7 to $26 million on conferences; executives dismissed in March 2016). https://www.cbsnews.com/news/cbs-news-investigates-wounded-warrior-project-spending/
491. GuideStar, Charity Navigator and BBB Wise Giving Alliance — *The Overhead Myth* (2013; overhead as a bad indicator). https://learn.guidestar.org/news/news-releases/2013/2013-06-17-overhead-myth
492. GiveWell — *How much does it cost to save a life?* (between ~$3,000 and $5,500 per life). https://www.givewell.org/how-much-does-it-cost-to-save-a-life
493. CharityWatch (Daniel Borochoff) — *Our charity rating process* (≥75% to program; ≤$25 to raise $100). https://www.charitywatch.org/our-charity-rating-process
494. CBS News / Patrick Radden Keefe (*Empire of Pain*) — *The Met removes the Sackler name* (philanthropy as "reputation laundering"). https://www.cbsnews.com/news/metropolitan-museum-of-art-sackler-family-name-removed-purdue-pharma/
495. Federal Trade Commission — *Reported fraud losses hit $12.5 billion in 2024* (fraud context). https://www.ftc.gov/news-events/news/press-releases/2025/03/new-ftc-data-show-big-jump-reported-losses-fraud-125-billion-2024
496. NPR — *Tesla ordered to pay $243 million in fatal Autopilot crash* (Benavides v. Tesla; $200 million in punitive damages; 33% liability). https://www.npr.org/2025/08/02/nx-s1-5490930/tesla-autopilot-crash-jury-240-million-florida

497. Electrek — *Tesla withheld data and misdirected police in Autopilot crash* (the collision snapshot at +3 min; the NAND image; the expert Alan Moore). https://electrek.co/2025/08/04/tesla-withheld-data-lied-misdirected-police-plaintiffs-avoid-blame-autopilot-crash/
498. National Transportation Safety Board — *Collision Between an Uber ATG Test Vehicle and a Pedestrian (HAR-19-03)* (5.6 s; emergency braking disabled). https://www.ntsb.gov/investigations/Pages/HWY18MH010.aspx
499. California DMV — *Statement on Cruise LLC suspension* (24-Oct-2023; the dragging footage, omitted). https://www.dmv.ca.gov/portal/news-and-media/dmv-statement-on-cruise-llc-suspension/
500. National Transportation Safety Board — *Cockpit Voice Recorders and Flight Data Recorders* (survivability; ≥88 parameters). https://www.ntsb.gov/news/Pages/cvr_fdr.aspx
501. Federal Register / FAA — *25-hour Cockpit Voice Recorder requirement for new aircraft* (2026; from 2 to 25 hours). https://www.federalregister.gov/documents/2026/02/02/2026-02110/25-hour-cockpit-voice-recorder-cvr-requirement-new-aircraft-production
502. eCFR — *49 CFR Part 563: Event Data Recorders* (pre-crash capture from 5 s to 20 s at 10 Hz). https://www.ecfr.gov/current/title-49/subtitle-B/chapter-V/part-563
503. DroneXL — *Drone pilot sentenced over Cape Canaveral Space Force base flights* (the geofencing logs contradict the operator). https://dronexl.co/2025/08/28/drone-pilot-sentenced-cape-canaveral-space-force-base/
504. CNN — *Jeju Air black boxes stopped recording before the crash* (the 4 decisive minutes, outside the box). https://www.cnn.com/2025/01/11/asia/south-korean-airline-black-boxes-stopped-recording-intl-hnk/index.html
505. National Transportation Safety Board — *EgyptAir 990 Aircraft Accident Brief (AAB-02/01)* (the same data, opposite causes). https://www.ntsb.gov/investigations/AccidentReports/Reports/AAB0201.pdf
506. Euronews — *Un fallo informático en las pulseras antimaltrato en España absuelve a decenas de agresores* (September 18, 2025; Telefónica → Vodafone+Securitas joint venture migration; data prior to March 20, 2024 inaccessible; State Prosecutor General's annual report; ≈4,500 active bracelets). https://es.euronews.com/2025/09/18/un-fallo-informatico-en-las-pulseras-antimaltrato-en-espana-absuelve-a-decenas-de-agresore
507. Diario Socialista — *Fallo en pulseras «antimaltrato» provoca absoluciones y pérdida de pruebas en casos de violencia machista* (September 18, 2025; "numerous acquittals and provisional dismissals"). https://diariosocialista.net/2025/09/18/fallo-en-pulseras-antimaltrato-provoca-absoluciones-y-perdida-de-pruebas-en-casos-de-violencia-machista/
508. OKDiario — *El Gobierno reconoce ahora que conocía desde «enero de 2024» los fallos en las pulseras antimaltrato*. https://okdiario.com/espana/pulseras-gobierno-reconoce-ahora-que-sabia-desde-enero-2024-fallos-pulseras-antimaltrato-15967835
509. La Nación (agencies) — *Redondo aborda en el Senado la gestión de pulseras para maltratadores tras ser reprobada en el Congreso* (September 2025; apologies to the victims; possible penalties for the contractor). https://www.lanacion.com.ar/agencias/redondo-aborda-este-miercoles-en-el-senado-la-gestion-de-pulseras-para-maltratadores-tras-ser-nid28092025/
510. Hay Derecho — Nuria Barona Martínez, *Pulseras antimaltrato: análisis jurídico y político de un sistema en crisis* (October 8, 2025). https://www.hayderecho.com/2025/10/08/pulseras-antimaltrato-analisis-juridico-politico/
511. Constitutional Court — *STC 114/1984, of November 29* (recording one's own conversation does not violate Art. 18.3 of the Constitution; admissible as documentary evidence; the intimate-sphere limit; the case's labor origin). https://hj.tribunalconstitucional.es/es/Resolucion/Show/367 · https://www.boe.es/buscar/doc.php?id=BOE-T-1984-27955
512. Legal Today — *La grabación de las conversaciones entre particulares y su manifestación jurisprudencial* (the Supreme Court's settled doctrine on participants' recordings); Iberley — *La validez de las grabaciones privadas aportadas como prueba en el proceso penal*. https://www.legaltoday.com/practica-juridica/derecho-penal/penal/la-grabacion-de-las-conversaciones-entre-particulares-y-su-manifestacion-jurisprudencial-2014-08-14/ · https://www.iberley.es/revista/la-validez-las-grabaciones-privadas-aportadas-como-prueba-proceso-penal-1221
513. Gerson Vidal (lawyers) — *¿Es delito grabar una conversación sin consentimiento?* (Art. 197.1: 1 to 4 years' prison and a 12-to-24-month fine; dissemination: 2 to 5 years); Noticias Jurídicas — *Voces ocultas: la delgada línea entre grabar legalmente y cometer un delito*. https://www.gersonvidal.com/blog/delito-grabar-conversacion/ ·

https://noticias.juridicas.com/conocimiento/tribunas/18542-voces-ocultas:-la-delgada-linea-entre-grabar-legalmente-y-cometer-un-delito-/

514. Digital Perito — *Clonación de voz: deepfakes de audio y detección forense* (convincing clones with 3-5 minutes of audio; some systems with 30 seconds; spectrogram and prosody analysis). https://digitalperito.es/glosario/clonacion-voz/
515. Economist & Jurist — *La prueba sintética en el proceso judicial: retos probatorios de los deepfakes* (challenging shifts the burden; the Supreme Court demands a technical, not generic, challenge; reinforced admission criteria); Rubén Blázquez (Instituto Vasco de Derecho Procesal) — *Deepfakes en el procedimiento probatorio*. https://www.economistjurist.es/noticias-juridicas/la-prueba-sintetica-en-el-proceso-judicial-retos-probatorios-de-los-deepfakes/ · https://www.institutovascodederechoprocesal.com/sites/default/files/files/Ruben%20Blazquez2.pdf
516. ICPM — *La Fiscalía señala que sólo el 0,01 % de las denuncias por violencia machista son falsas* (79 final convictions, 0.0075%). https://www.icpm.es/actualidad_noticias.asp?id=2382
517. Efeminista — *Casi 3 millones de denuncias por violencia de género en España desde 2004 y solo un 0,0082 % son falsas* (2009-2024 series); ANAVID — *¿De dónde sale la cifra del 0,01 %?* (methodological nuance: only falsehood attested with intent). https://efeminista.com/denuncias-falsas-violencia-de-genero-espana/ · https://anavid.es/de-donde-sale-la-cifra-del-001-de-denuncias-falsas-de-violencia-de-genero-por-parte-de-la-fiscalia-del-estado-y-el-cgpj/
518. Motorpasión — *¿Puedes usar una dashcam en España?* (AEPD criteria: recording limited to the incident, ~20 seconds before and after, preferably without sound; valid as evidence via legitimate interest and judicial protection); Adasfera — *Legalidad dashcam España: guía DGT y AEPD (2026)*. https://www.motorpasion.com/seguridad/pregunta-puedes-usar-dashcam-espana-respuesta-puedes-meterte-buen-lio-no-sabes-como · https://adasfera.com/2026/01/legalidad-dashcam-espana-dgt-multasson-legales-las-dashcams-en-espana-en-2026-guia-definitiva-dgt-y-aepd/
519. TodoDashcam — *Legislación sobre dashcams en España y Europa* (Portugal: Law 58/2019, banned even to carry one; Germany: admission limited to the incident; France: CNIL, driver responsible for the processing). https://www.tododashcam.es/pages/blog-legislacion-dashcams-espana
520. AlertCops (State Secretariat for Security, Ministry of the Interior) — *AlertCops incorpora un botón SOS* (5 presses in under 6 seconds; location + 10 seconds of audio to the security forces; AVDG code for victims of gender violence). https://alertcops.ses.mir.es/publico/alertcops/actualidad/detalleActualidad.html?entrada=AlertCops-incorpora-un-bot-n-SOS
521. Rastreator — *Guardia Civil y UNESPA, contra el fraude* (≈400,000 fraud attempts a year; ≈€2 billion; 6 in 100 injuries; staged crashes: 3.2% of those declared, 18% of the investigated cost; GC-UNESPA Operating Procedure). https://www.rastreator.com/seguros/noticias/guardia-civil-y-unespa-luchan-contra-el-fraude-en-el-seguro.aspx
522. Infobae — *Detenidas 32 personas que fingían accidentes de tráfico para estafar al seguro* (July 2, 2025; seven staged crashes in eight months). https://www.infobae.com/espana/2025/07/02/detenidas-32-personas-que-fingian-accidentes-de-trafico-para-estafar-al-seguro-una-agencia-de-detectives-presento-la-denuncia/
523. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; the content never leaves the device; voluntary disclosure). See chapters 13 and 15.
524. TechCrunch — *23andMe confirms hackers stole ancestry data on 6.9 million users* (2023; the lists segmented by ethnic origin). https://techcrunch.com/2023/12/04/23andme-confirms-hackers-stole-ancestry-data-on-6-9-million-users/
525. Yaniv Erlich, Tal Shor, Itsik Pe'er and Shai Carmi (Science) — *Identity inference of genomic data using long-range familial searches* (2018; 2% of the population is enough). https://www.science.org/doi/10.1126/science.aau4832
526. CNBC — *23andMe bankruptcy: selling and deleting DNA* (2025; +526% traffic; ~1.9 million deletion requests). https://www.cnbc.com/2025/03/30/23andme-bankruptcy-selling-deleting-dna-genetic-testing.html
527. New York Attorney General (Letitia James) — *AG James sues 23andMe to protect genetic data* (a coalition of 28 attorneys general). https://ag.ny.gov/press-release/2025/attorney-general-james-sues-23andme-protect-new-yorkers-genetic-data
528. Bio-IT World — *Anne Wojcicki buys back 23andMe for $305M* (Regeneron, $256 million → TTAM, $305 million; closing on 14-Jul-2025). https://www.bio-itworld.com/news/2025/07/02/anne-wojcicki-buys-back-23andme-for-305m-promises-data-security

529. Information Commissioner's Office (United Kingdom) — *ICO fines 23andMe £2.31 million* (155,592 UK residents; the penalty punishes the custody). https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2025/06/23andme-fined-for-failing-to-protect-uk-users-genetic-data/
530. NPR — *In hunt for the Golden State Killer, investigators uploaded his DNA to a genealogy site* (the DNA was uploaded by a relative). https://www.npr.org/sections/thetwo-way/2018/04/27/606624218/in-hunt-for-golden-state-killer-investigators-uploaded-his-dna-to-genealogy-site
531. The Legal Genealogist — *GEDmatch reverses course* (opt-out by default; only ~14% opted in). https://www.legalgenealogist.com/2019/05/19/gedmatch-reverses-course/
532. National Human Genome Research Institute — *GINA: Genetic Information Nondiscrimination Act (2008)* (does not cover life, disability or long-term care insurance). https://www.genome.gov/27026050/president-bush-signs-the-genetic-information-nondiscrimination-act-of-2008
533. Taner Kuru (Computer Law & Security Review) — *Investigative genetic genealogy in Europe* (2025; uploading your DNA does not make your relatives' "public"). https://www.sciencedirect.com/science/article/pii/S0267364925000019
534. Fortune — *Tesla lawyers argue recordings of Musk touting Autopilot could be deepfakes* (Huang v. Tesla; Judge Evette Pennypacker). https://fortune.com/2023/04/27/elon-musk-lawyers-argue-recordings-of-him-touting-tesla-autopilot-safety-could-be-deepfakes/
535. 9to5Mac — *Rittenhouse lawyer's pinch-to-zoom argument* (the burden shifted onto the prosecution; 20 minutes for an Apple expert). https://9to5mac.com/2021/11/11/rittenhouse-lawyer-pinch-to-zoom/
536. NBC News — *Washington state judge blocks use of AI-enhanced video evidence* (State v. Puloka; Judge Leroy McCullough; 29-Mar-2024). https://www.nbcnews.com/news/us-news/washington-state-judge-blocks-use-ai-enhanced-video-evidence-rcna141932
537. NPR — *Baltimore police caught planting drugs in body-cam footage* (Officer Pinheiro; the 30-second soundless buffer). https://www.npr.org/sections/thetwo-way/2017/07/20/538279258/baltimore-police-caught-planting-drugs-in-body-cam-footage-public-defender-says
538. ProPublica — *How police have undermined the promise of body cameras* (only 42% of 79 killings; a false "does not exist" in 1 in 5). https://www.propublica.org/article/how-police-undermined-promise-body-cameras
539. Robert Chesney and Danielle Keats Citron — *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security* (107 California Law Review 1753, 2019; the "liar's dividend"). https://www.californialawreview.org/print/deep-fakes-a-looming-challenge-for-privacy-democracy-and-national-security
540. U.S. Courts — *Advisory Committee on Evidence Rules: proposed Rule 707 (machine-generated evidence) and 901(c)* (comment period until 16-Feb-2026). https://www.uscourts.gov/sites/default/files/document/2025-11_evidence_rules_commitee_agenda_book_final.pdf
541. American Bar Association — *New rules for self-authenticating electronic evidence (FRE 902(13)/(14))* (the hash value self-authenticates digital copies). https://www.americanbar.org/groups/litigation/resources/newsletters/trial-evidence/new-rules-self-authenticating-electronic-evidence/
542. CYFOR / Expatriate Law — *Deepfake audio evidence used in a UK court* (audio edited after the call to insert threats). https://cyfor.co.uk/deepfake-audio-evidence-used-in-uk-court-to-discredit-father/
543. BBC (via LBC) — *150 reports of police officers misusing body cameras* (England and Wales; Jim Colwell, NPCC). https://www.lbc.co.uk/article/150-reports-police-officers-misusing-body-cameras-investigation-england-wales-DWzgZZ_2/
544. Enigma Forensics — *Judge sanctions hospital for not producing complete EMR audit trail* (Prieto v. Rush University Medical Center; Judge James O'Hara; "information blocking"). https://enigmaforensics.com/blog/judges-order-sanctions-hospital-not-producing-complete-emr-audit-trail/
545. Cornell Law (Legal Information Institute) — *45 CFR §164.312: HIPAA technical safeguards* (mandatory audit controls; ePHI integrity). https://www.law.cornell.edu/cfr/text/45/164.312
546. Cornell Law (Legal Information Institute) — *21 CFR §11.10: FDA electronic records* (timestamped audit trails; a change cannot obscure the previous record). https://www.law.cornell.edu/cfr/text/21/11.10

547. HHS Office of Inspector General — *Not All Recommended Fraud Safeguards Have Been Implemented in Hospital EHR Technology* (OEI-01-11-00570, 2013; only ~25% with copy-paste policies). https://oig.hhs.gov/oei/reports/oei-01-11-00570.asp
548. HHS Office of Inspector General — *Trend Toward More Expensive Inpatient Hospital Stays in Medicare* (OEI-02-18-00380, 2021; almost half the spending at the highest severity level). https://oig.hhs.gov/oei/reports/OEI-02-18-00380.pdf
549. ProPublica / HHS OCR — *UCLA Health System pays $865,500 to settle celebrity-privacy allegations* (2011; the snooping in the charts). https://www.propublica.org/article/ucla-health-system-pays-865000-to-settle-celebrity-privacy-allegations
550. U.S. Department of Justice — *Former UCLA healthcare worker sentenced (Huping Zhou)* (323 accesses; the first prison sentence for a HIPAA violation). https://www.justice.gov/archive/usao/cac/Pressroom/pr2010/004.html
551. Federal Register / ONC — *21st Century Cures Act: interoperability and information blocking* (2020; patient access to their health information). https://www.federalregister.gov/documents/2020/05/01/2020-07419/21st-century-cures-act-interoperability-information-blocking-and-the-onc-health-it-certification
552. Kantara Initiative — *Consent Receipt Specification v1.1* (basis of the ISO/IEC TS 27560:2023 standard). https://kantara.atlassian.net/wiki/spaces/archive/pages/3508790/Consent+Receipt+Specification
553. Dean F. Sittig and Adam Wright (Applied Clinical Informatics) — *Identifying a Clinical Informatics or EHR Expert Witness* (2023; the log is evidence that must be interpreted). https://pmc.ncbi.nlm.nih.gov/articles/PMC10033222/
554. Stanford Internet Observatory / 404 Media — *Largest Dataset Powering AI Images Removed After Discovery of Child Sexual Abuse Material* (December 2023; 3,226 suspicious cases in LAION-5B; recommendation to deprecate the models). https://www.404media.co/laion-datasets-removed-stanford-csam-child-abuse/
555. LAION — *Releasing Re-LAION-5B: transparent iteration on LAION-5B with additional safety fixes* (August 2024; 2,236 links removed). https://laion.ai/blog/relaion-5b/
556. JFrog Security Research — *Data Scientists Targeted by Malicious Hugging Face ML Models with Silent Backdoor* (2024; the *nullifAI* technique, *pickle* abuse, backdoor). https://jfrog.com/blog/data-scientists-targeted-by-malicious-hugging-face-ml-models-with-silent-backdoor/
557. The Hacker News — *Over 100 Malicious AI/ML Models Found on Hugging Face Platform* (March 2024). https://thehackernews.com/2024/03/over-100-malicious-aiml-models-found-on.html
558. EU Artificial Intelligence Act — *Article 53: Obligations for providers of general-purpose AI models* (Annex XI technical documentation; training content summary; penalties of up to €15 million or 3%). https://artificialintelligenceact.eu/article/53/
559. European Commission — *General-Purpose AI Models in the AI Act: Questions & Answers* (the training summary template; application from August 2, 2025; Article 50 marking from August 2, 2026). https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers
560. CycloneDX (OWASP) — *Machine Learning Bill of Materials (AI/ML-BOM)* (datasets, models and configurations with provenance; since v1.5, 2023). https://cyclonedx.org/capabilities/mlbom/
561. *Operationalising AI bills of materials for verifiable AI provenance and lifecycle assurance* (Frontiers in Computer Science, 2026). https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2026.1735919/full
562. Margaret Mitchell et al. — *Model Cards for Model Reporting* (2019). https://arxiv.org/abs/1810.03993
563. Google / C2PA — *How Google and the C2PA are increasing transparency for gen AI content* (SynthID: more than 100 billion pieces of content marked according to Google I/O 2026 (May 2026); content credentials). https://blog.google/innovation-and-ai/products/google-gen-ai-content-transparency-c2pa/
564. Cathy O'Neil — *Armas de destrucción matemática* (Capitán Swing, 2018).
565. Council of the EU / Pinsent Masons (2026). https://www.pinsentmasons.com/out-law/news/rules-high-risk-ai-delayed-under-eu-omnibus-deal
566. Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 6.
567. VentureBeat (2020). https://venturebeat.com/ai/amsterdam-and-helsinki-launch-algorithm-registries-to-bring-transparency-to-public-deployments-of-ai

568. CBC News — *Air Canada found liable for chatbot's bad advice on bereavement rates* (Moffatt v. Air Canada, B.C. Civil Resolution Tribunal, February 2024: the chatbot invented a bereavement fare policy; CA$812 in compensation on 1,630.36 paid; "Air Canada did not take reasonable care to ensure its chatbot was accurate"; "it is responsible for all the information on its website — whether it comes from a static page or a chatbot"); Forbes — *What Air Canada Lost In 'Remarkable' Lying AI Chatbot Case* (the claim that the chatbot was "a separate legal entity"); McCarthy Tétrault — *Moffatt v. Air Canada: A Misrepresentation by an AI Chatbot*. https://www.cbc.ca/news/canada/british-columbia/air-canada-chatbot-lawsuit-1.7116416 · https://www.forbes.com/sites/marisagarcia/2024/02/19/what-air-canada-lost-in-remarkable-lying-ai-chatbot-case/ · https://www.mccarthy.ca/en/insights/blogs/techlex/moffatt-v-air-canada-misrepresentation-ai-chatbot
569. Visa — *Intelligent Commerce* (limited-scope tokenized credentials for agents; MCP server; a $200 token cannot authorize $500); Mastercard — *Agent Pay* (announced 29-4-2025: Agentic Tokens binding credential, agent, merchant scope and consent policy; "Agent Pay for Machines," June 2026); Google AP2 (September 2025, 60+ partners: Intent Mandate + Cart Mandate signed as verifiable credentials); Eco/Paz — 2026 agentic commerce guides (runtimes: OpenAI Operator, Google Project Mariner, Anthropic Computer Use, Microsoft Copilot). https://www.visa.com/en-us/solutions/intelligent-commerce · https://www.mastercard.com/us/en/news-and-trends/press/2026/june/mastercard-launches-agent-pay-for-machines.html · https://eco.com/support/en/articles/14839400-what-is-agentic-commerce-the-2026-guide · https://www.paz.ai/glossary/agentic-commerce-protocol-acp
570. t2p — product manual (certify any file — the mandate and the logs, only the fingerprint travels; web page capture with full image + source code; signature with video selfie, two random liveness challenges and the device's biometric verification; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the models' provenance, in chapter 79; public sector agents, in chapter 107.
571. Voicebot.ai — *Tom Hanks Warns of Deepfake Hanks Dental Ad* (October 2023; "I have nothing to do with it"). https://voicebot.ai/2023/10/02/tom-hanks-warns-of-deepfake-hanks-dental-ad/
572. NPR — *OpenAI pulls AI voice that was compared to Scarlett Johansson in the movie 'Her'* (May 2024; forensic analysis by Arizona State University: more similar than 98% of actresses). https://www.npr.org/2024/05/20/1252495087/openai-pulls-ai-voice-that-was-compared-to-scarlett-johansson-in-the-movie-her
573. SAG-AFTRA / Perkins Coie — *How the 2023 SAG-AFTRA contract addresses generative AI and digital replicas* (informed consent; specific description of the use; new consent for reuse). https://perkinscoie.com/insights/blog/generative-ai-movies-and-tv-how-2023-sag-aftra-and-wga-contracts-address-generative
574. Davis Wright Tremaine — *Tennessee Passes ELVIS Act (Ensuring Likeness, Voice and Image Security Act)* (in force July 1, 2024; the voice within the right of publicity; AI tool liability). https://www.dwt.com/blogs/artificial-intelligence-law-advisor/2024/04/tennessee-elvis-act-ai-voice-replica
575. United States Senate Judiciary Committee — *NO FAKES Act (Nurture Originals, Foster Art, and Keep Entertainment Safe)* (a federal right over the digital replica; revised 2026 version, approved by the Senate Judiciary Committee on June 18, 2026). https://www.judiciary.senate.gov/no-fakes-act
576. World Economic Forum — *Deepfake legislation: Denmark takes action* (introduced June 2025; a copyright-style right over face, body and voice; removal and compensation; parody exceptions; in parliamentary procedure, entry into force expected mid-2026). https://www.weforum.org/stories/2025/07/deepfake-legislation-denmark-digital-id/
577. Variety — *Scarlett Johansson Reacts to OpenAI Sky Voice: 'Shocked' and 'Angered'* (May 2024). https://variety.com/2024/digital/news/scarlett-johansson-responds-shocked-angered-openai-chatgpt-her-1236011135/
578. SAG-AFTRA — *Digital Replicas 101* (typology and the performer's consent rights). https://www.sagaftra.org/sites/default/files/sa_documents/DigitalReplicas.pdf
579. Wikipedia — *ELVIS Act* (Ensuring Likeness, Voice and Image Security Act of 2024; reform of the Personal Rights Protection Act of 1984). https://en.wikipedia.org/wiki/ELVIS_Act
580. Marc Goodman — *Delitos del futuro* (Ariel, 2015).
581. Wikipedia, *Jason Palmer (politician)*; TIME, *Albania's AI minister Diella* (2025). https://time.com/7324934/albania-ai-minister-diella/

582. Légifrance — *Law No. 2024-120, of February 19, 2024, guaranteeing respect for children's image rights* (amendment of Art. 371-1 of the Civil Code: parental authority includes protecting the minor's private life and image right; protection exercised jointly by both parents, involving the child according to age and maturity; judicial delegation of the exercise of the image right in case of grave harm to dignity or moral integrity); DAJ/economie.gouv.fr — *Respect du droit à l'image des enfants: les apports de la loi* (context: a child appears on average in 1,300 photographs online before age 13; risks: identity theft, blackmail, cyberbullying); Actu-Juridique — *Les 5 apports de la loi*. https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000049163317 · https://www.economie.gouv.fr/daj/lettre-de-la-daj-respect-du-droit-limage-des-enfants-les-apports-de-la-loi-du-19-fevrier-2024 · https://www.actu-juridique.fr/civil/personnes-famille/le-respect-du-droit-a-limage-des-enfants-et-les-5-apports-de-la-loi-du-19-fevrier-2024/
583. The Spanish framework (LO 1/1982 on honor, privacy and one's own image; LO 1/1996 on the legal protection of minors; intervention of the public prosecutor; case law requiring the agreement of both parents to publish images of the children). See chapter 81 for general image rights and chapter 53 for co-parenting records.
584. t2p — product manual (document scanner; certified screenshot and email; certify any file — only the fingerprint travels, the content never leaves the device; signature with video selfie; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; co-parenting, in chapter 53; intimate dissemination, in chapter 83.
585. Euronews — *The Almendralejo case: when AI deepfakes are used to undress teenagers* (September 2023; a mother's testimony; victims aged 12-14). https://www.euronews.com/next/2023/11/15/the-case-of-almendralejo-when-deepfakes-are-used-to-undress-teenagers-euronews-tech-talks-
586. Scottish Legal News / Reuters — *Spain: court punishes schoolboys for spreading AI deepfakes of girls* (July 2024; 15 minors convicted; 20 offenses of child abuse material and 20 against moral integrity). https://www.scottishlegal.com/articles/spain-court-punishes-schoolboys-for-spreading-ai-deepfakes-of-girls
587. Wikipedia — *Taylor Swift deepfake pornography controversy* (January 2024; one post viewed more than 47 million times; X blocks searches of her name). https://en.wikipedia.org/wiki/Taylor_Swift_deepfake_pornography_controversy
588. Congress.gov / The White House — *TAKE IT DOWN Act* (S.146; signed May 19, 2025; takedown within 48 h; platform compliance enforced by the FTC since May 19, 2026; first criminal conviction in April 2026). https://www.congress.gov/bill/119th-congress/senate-bill/146
589. Council of Europe — *Directive (EU) 2024/1385 of 14 May 2024 on combating violence against women and domestic violence* (in force June 13, 2024; criminalizes manipulated intimate material; takedown orders and preservation of evidence). https://www.coe.int/en/web/cyberviolence/-/european-parliament-and-council-directive-eu-2024-1385-of-14-may-2024-on-combating-violence-against-women-and-domestic-violence
590. Euronews — *Nude deepfakes: is the EU doing enough to tackle the issue?* (2024; more than 90% of deepfake content is pornographic, almost all targeting women). https://www.euronews.com/next/2024/03/17/nude-deepfakes-is-the-eu-doing-enough-to-tackle-the-issue
591. StopNCII.org — *How StopNCII.org works* (fingerprint generated on the device; the image never leaves; more than 500,000 images protected; partners Meta, TikTok, Reddit, Google). https://stopncii.org/how-it-works/
592. Skadden — *'Take It Down Act' requires online platforms to remove unauthorized intimate images and deepfakes when notified* (analysis of the notice-and-takedown mechanism). https://www.skadden.com/insights/publications/2025/06/take-it-down-act
593. RAINN — *TAKE IT DOWN Act* (the perspective of victims of sexual violence). https://rainn.org/federal-legislation/take-it-down-act/
594. The Markup (2024). https://themarkup.org/artificial-intelligence/2024/12/11/1-in-6-congresswomen-targeted-by-ai-generated-sexually-explicit-deepfakes
595. The Local (Italy) — *Italy unanimously passes anti-cyberbullying law* (May 2017; Law 71/2017; Carolina Picchio, 14, in 2013; Senator Elena Ferrara). https://www.thelocal.it/20170517/italy-unanimously-passes-anti-cyberbullying-law-online-abuse-trolling
596. Pew Research Center — *Teens and Cyberbullying 2022* (December 15, 2022; 46%; name-calling 32%, rumors 22%, threats 10%). https://www.pewresearch.org/internet/2022/12/15/teens-and-cyberbullying-2022/

597. WHO-Europe and the HBSC study — *One in six school-aged children experiences cyberbullying* (March 27, 2024; 16%; 44 countries; ≈280,000 young people). https://www.who.int/europe/news/item/27-03-2024-one-in-six-school-aged-children-experiences-cyberbullying--finds-new-who-europe-study
598. ANAR Foundation and Mutua Madrileña Foundation — *VIII Informe de Prevención del Acoso Escolar* (2025; 12.3% vs 9.4%; 14.2% with AI; 91.7% from the same school; 17,562 students). https://www.anar.org/crece-el-acoso-escolar-impulsado-por-el-ciberbullying-y-el-uso-de-la-ia-para-cometerlo/
599. ADL — *Online Hate and Harassment: The American Experience 2024* (June 2024; teenagers 50%; 36% spilling into the physical world). https://www.adl.org/resources/report/online-hate-and-harassment-american-experience-2024
600. CBS News / ITV News — *Molly Russell inquest* (September 30, 2022; coroner Andrew Walker; 2,100 of 16,300 posts; the world's first such ruling). https://www.cbsnews.com/news/molly-russell-suicide-death-coroner-social-media-instagram-pinterest/
601. Candice Odgers — *The great rewiring: is social media really behind an epidemic of teenage mental illness?* (Nature, March 2024; a reply to Jonathan Haidt, *The Anxious Generation*, 2024). https://www.nature.com/articles/d41586-024-00902-2
602. CNN — *Pikesville principal's voice on racist recording was an AI-generated deepfake* (Dazhon Darien / Eric Eiswert; January-April 2024). https://www.cnn.com/2024/04/26/us/pikesville-principal-maryland-deepfake-cec/index.html
603. PNAS — *X under Musk's leadership: substantial hate and no reduction in inauthentic activity* (2025; +≈50% hate; +70% "likes"; coordinated activity not reduced). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11819532/
604. European Commission — *Guidelines on the protection of minors under the Digital Services Act* (July 14, 2025; cyberbullying, age verification; systemic risks = illegal hate speech). https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-protection-minors
605. La Moncloa — *Anteproyecto de Ley Orgánica para la protección de las personas menores de edad en los entornos digitales* (Council of Ministers, June 4, 2024; "online restraining order"; age verification). https://www.lamoncloa.gob.es/consejodeministros/resumenes/Paginas/2024/040624-rueda-de-prensa-ministros.aspx
606. Criminal Code (Spain) — Arts. 510 (hate crime; 1-4 years' prison) and 172 ter (harassment/stalking; aggravated by the victim's age). https://www.boe.es/buscar/act.php?id=BOE-A-1995-25444
607. Ofcom / GOV.UK — *Online Safety Act 2023: protection of children codes* (duties in force July 25, 2025; bullying and hateful content; fines up to 10% of worldwide turnover). https://www.gov.uk/government/collections/online-safety-act
608. eSafety Commissioner (Australia) — *Social media age restrictions* (Act 2024; under-16s from December 10, 2025; the first country). https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions
609. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container). See chapters 13 and 15.
610. INCIBE / Infobae — *El teléfono de tu hijo no se dañó: la estafa en WhatsApp que pide dinero para una emergencia* (May 8, 2025; the mechanics of the "hi Mom" scam, variants, verification advice); National Police (alert) — OKDiario; Valencia News — *Así funciona la estafa del «hijo en apuros» por SMS y WhatsApp*. https://www.infobae.com/tecno/2025/05/08/el-telefono-de-tu-hijo-no-se-dano-la-estafa-en-whatsapp-que-pide-dinero-para-solucionar-una-emergencia/ · https://okdiario.com/curiosidades/policia-nacional-alerta-sobre-esta-estafa-hola-mama-telefono-esta-roto-10582931 · https://valencianews.es/tendencias/mama-se-me-ha-roto-el-movil-asi-funciona-estafa-del-hijo-en-apuros-sms-y-whatsapp/
611. HousingWire — *FBI: Seniors lost $7.75B to cybercrime in 2025 — a 59% jump* (201,266 complaints from over-60s; $4.9 billion in 2024, +43%; >3,100 senior complaints with AI and >$352 million; >$5 million in distress scams with cloned voices); PPC Land — *AI-linked fraud costs Americans $893 million* (22,364 AI complaints in 2025). https://www.housingwire.com/articles/fbi-seniors-cybercrime-2025/ · https://ppc.land/ai-linked-fraud-costs-americans-893-million-fbi-reports-for-2025/
612. FBI / IC3 — *Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud* (PSA-241203, December 3, 2024; voice cloning in distress scams and family impersonation). https://www.ic3.gov/PSA/2024/PSA241203

613. Factchequeado — *Qué son las estafas de romance y cómo identificarlas* (FTC: ≈60,000 complaints and $1.2 billion in 2024; the 65-74 bracket with the highest losses); AARP — *Los estafadores de romance engañan a sus víctimas con fraudes de criptomonedas.* https://factchequeado.com/teexplicamos/20240126/estafa-romance-scam/ · https://www.aarp.org/espanol/dinero/estafas-y-fraudes/info-2023/enganos-de-romance-con-criptomoneda.html
614. Wikipedia — *Estafa romántica* (the Morata de Tajuña crime, 2023: three elderly siblings murdered; a scam of up to €263,000). https://es.wikipedia.org/wiki/Estafa_rom%C3%A1ntica
615. Ministry of the Interior — *Detenidas dos personas por la estafa conocida como «Romance Scam», timo del amor*. https://www.interior.gob.es/opencms/es/detalle/articulo/Detenidas-dos-persona-por-la-estafa-conocida-como-Romance-Scam-timo-del-amor/
616. Infobae — *Francia anuncia control de residencias de ancianos tras escándalo* (February 1, 2022; Victor Castanet's *Les Fossoyeurs*; Orpea ≈260,000 residents; rationing; criminal investigations into mistreatment, misappropriation and labor violations; involuntary manslaughter complaints); infoLibre — *El último escándalo de Orpea* (state injection of €605 million for a majority stake). https://www.infobae.com/america/agencias/2022/02/01/francia-anuncia-control-de-residencias-de-ancianos-tras-escandalo/ · https://www.infolibre.es/politica/ultimo-escandalo-orpea-negocios-trama-societaria-oculta-luxemburgo_1_1232899.html
617. elDiario.es — *Condenada la trabajadora de una residencia cazada con una cámara oculta mientras maltrataba a una interna* (El Escorial; a €480 fine and €3,000 compensation upheld by the Provincial Court of Madrid; the daughters' camera endorsed as evidence). https://www.eldiario.es/madrid/condenada-trabajadora-residencia-cazada-camara-oculta-maltrataba-interna_1_12498186.html
618. El Foro de Labos — *El Tribunal Supremo confirma la validez de una prueba obtenida con cámara oculta* (2022); Pardo Geijo Abogados — *¿Es legal poner cámaras ocultas en casa?* (well-founded suspicion of a serious crime; temporary, focused camera; triple proportionality test). https://www.elforodelabos.es/2022/09/el-tribunal-supremo-confirma-la-validez-de-una-prueba-obtenida-con-camara-oculta/ · https://pardogeijo.com/es-legal-grabar-camaras-ocultas-casa/
619. BOE — *Law 8/2021, of June 2, reforming civil and procedural law to support persons with disabilities in the exercise of their legal capacity* (support, not substitution; voluntary measures before a notary; capacity to make a will; rendering of accounts; exceptional guardianship); El Notario — *Una visión práctica de la Ley 8/2021*. https://www.boe.es/buscar/act.php?id=BOE-A-2021-9233 · https://www.elnotario.es/practica-juridica/11767-una-vision-practica-de-la-ley-8-2021-de-2-de-junio-sobre-personas-con-discapacidad
620. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; proof of life with challenge; content never off the device; voluntary disclosure). See chapters 13, 15 and 64.
621. Atlantic Council — *Russian War Report: Hacked news program and deepfake video spread false Zelenskyy claims* (March 16, 2022; hacking of Ukraine 24's ticker and website; surrender deepfake; "childish provocation" selfie refutation). https://www.atlanticcouncil.org/blogs/new-atlanticist/russian-war-report-hacked-news-program-and-deepfake-video-spread-false-zelenskyy-claims/
622. NPR — *Impostor uses AI to impersonate Rubio and contact foreign and U.S. officials* (July 2025; the secretary of state's cloned voice on Signal; State Department cable of July 3). https://www.npr.org/2025/07/09/nx-s1-5462195/impostor-ai-impersonate-rubio-foreign-officials
623. Clarity — *Iranian Deepfake of Israeli Defense Minister Airs Live* (April 2025; Channel 14's unverified rebroadcast of a Yoav Gallant deepfake; live cut and apology). https://www.getclarity.ai/ai-deepfake-blog/iranian-deepfake-of-israeli-defense-minister-airs-live---and-how-clarity-caught-it-instantly
624. Romania Insider — *Romanian central bank governor Mugur Isarescu targeted by deepfake* (February 2024; a fraudulent-investment deepfake; a pattern across several central banks). https://www.romania-insider.com/bnr-mugur-isarescu-deepfake-february-2024
625. Wikipedia — *2018 Hawaii false missile alert* (January 13, 2018; false WEA ballistic missile alert; 38-minute correction; genuine error). https://en.wikipedia.org/wiki/2018_Hawaii_false_missile_alert
626. FEMA — *Integrated Public Alert & Warning System (IPAWS)* (official alerts digitally signed by the issuing authority and authenticated by IPAWS-OPEN; the baseline of authentication at emission). https://www.fema.gov/emergency-managers/practitioners/integrated-public-alert-warning-system

627. Singapore Police Force — *Footage from Zoom video conference involving impersonation of senior government officials* (May 16, 2026; deepfake of PM Lawrence Wong, President Tharman and MAS representatives; a 4.9 million Singapore dollar loss by a single victim). https://www.police.gov.sg/media-hub/news/2026/05/20260516_footage_from_zoom_video_conference_involving_impersonation_of_senior_govern
628. Bulletin of the Atomic Scientists — *AI misinformation is threatening emergency communications. Here's how to fix that* (September 2025; false official alerts, the "liar's dividend" and the authentication of emergency communications). https://thebulletin.org/2025/09/ai-misinformation-is-threatening-emergency-communications-heres-how-to-fix-that/
629. Fortune — *A deepfake 'CFO' tricked British design firm Arup in $25 million fraud* (Arup confirmed as the victim; HK$200M/US$25.6M; 15 transfers; CIO Rob Greig's statement). https://fortune.com/europe/2024/05/17/arup-deepfake-fraud-scam-victim-hong-kong-25-million-cfo/
630. FinCEN — *FinCEN Issues Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions* (alert FIN-2024-Alert004, November 13, 2024; recommends live verification and phishing-resistant MFA). https://www.fincen.gov/news/news-releases/fincen-issues-alert-fraud-schemes-involving-deepfake-media-targeting-financial
631. FBI/IC3 — *Senior US Officials Impersonated in Malicious Messaging Campaign* (PSA250515, May 15, 2025; AI-generated voice/text; out-of-band verification and pre-agreed phrases). https://www.ic3.gov/PSA/2025/PSA250515
632. Fortune — *Ferrari exec foils deepfake attempt by asking the scammer a question only CEO Benedetto Vigna could answer* (July 2024; a voice clone foiled with a control question). https://fortune.com/2024/07/27/ferrari-deepfake-attempt-scammer-security-question-ceo-benedetto-vigna-cybersecurity-ai/
633. Marketing-Interactive — *WPP CEO Mark Read targeted by deepfake AI scam* (May 2024; fake WhatsApp, Teams, cloned voice and YouTube footage; failed attempt; Read's warning). https://www.marketing-interactive.com/wpp-ceo-mark-read-deepfake-ai-scam
634. Tookitaki — *The Deepfake Deception: How a CEO Video Call Cost $499K in Singapore* (March 2025; a finance director authorizes $499,000 on a Zoom call with every participant a deepfake). https://www.tookitaki.com/blog/deepfake-ceo-scam-singapore-2025
635. Nacha — *FBI's IC3 Finds Almost $8.5 Billion Lost to Business Email Compromise in Last Three Years* (~$3 billion/year; ~$8.5 billion over three years; the payment fraud baseline). https://www.nacha.org/news/fbis-ic3-finds-almost-85-billion-lost-business-email-compromise-last-three-years
636. Trend Micro — *Deepfake CFO Video Calls Result in $25MM in Damages* (technical analysis of the multi-participant video call attack pattern used against Arup). https://www.trendmicro.com/en_us/research/24/b/deepfake-video-calls.html
637. World Economic Forum — *Cybercrime: Lessons learned from a $25m deepfake attack* (February 2025; analysis of the Arup case and the control failures; the out-of-band verification argument). https://www.weforum.org/stories/2025/02/deepfake-ai-cybercrime-arup/
638. 9to5Mac (on WIRED's reporting) — *Latest romance scam tactic uses two phones with face-swap app* (Apr-2024: the "Yahoo Boys" and real-time face-swap inside video calls — two phones on stands and a ring light, or a laptop with a webcam; they show off the setups on Telegram and sell them; Rachel Tobac: the apps have improved dramatically in a year). https://9to5mac.com/2024/04/18/romance-scam-face-swap/
639. iProov — *Threat Intelligence Report 2025: Remote Identity Under Attack* (native virtual camera attacks +2,665%; face swaps +300% versus 2023, aimed at systems with liveness detection protocols; ~24,000 users selling attack technologies as a service; image-to-video conversion as the new vector; 127 face-swap tools, 91 virtual cameras and 10 emulators catalogued). https://www.iproov.com/reports/threat-intelligence-report-2025-remote-identity-attack · https://www.iproov.com/press/annual-identity-verification-threat-intelligence-report
640. Gartner (via FutureCIO) — *Deepfakes to render identity verification and authentication solutions inefficient* (Feb-2024 prediction: by 2026, 30% of companies will no longer consider identity verification and authentication solutions reliable in isolation, due to deepfakes against facial biometrics; injection attacks +200% in 2023; presentation attack detection standards do not cover digital injection; Akif Khan: combine PAD, IAD and image inspection). https://futurecio.tech/deepfakes-to-render-identity-verification-and-authentication-solutions-inefficient/

641. FBI / IC3 — *Deepfakes and Stolen PII Utilized to Apply for Remote Work Positions* (PSA of 28-Jun-2022: deepfakes and stolen personal data used to apply for remote IT, programming, database and software jobs — some with access to customer data and corporate systems; signals: lips and sound out of sync, coughs or sneezes that do not match the image). https://www.ic3.gov/PSA/2022/psa220628
642. Decrypt — *People Are Using Basic AI to Bypass KYC — But Should You?* (Feb-2024: OnlyFake, fake IDs and passports generated "with neural networks" for $15 — U.S., Italy, China, Russia, Argentina, Czechia, Canada; 404 Media passed the KYC of the exchange OKX, which uses the third-party verifier Jumio; users used them to open bank accounts and unlock crypto accounts). https://decrypt.co/216188/ai-generated-fake-id-bypass-kyc-aml-banks-crypto-onlyfakes
643. Adler, Hitzig, Jain et al. (arXiv) / MIT News — *Personhood credentials: Artificial intelligence and the value of privacy-preserving tools to distinguish who is real online* (Aug-2024; 32 authors, with researchers from MIT, OpenAI and Microsoft: digital credentials to prove you are a real person — not an AI — without revealing any personal data; in-person or institutional issuance; cryptographic verification; "AI is moving very fast, certainly much faster than the speed at which governments adapt"). https://arxiv.org/abs/2408.07892 · https://news.mit.edu/2024/3-questions-proving-humanity-online-0816
644. European Commission / Corbado — the European Digital Identity Regulation (EU) 2024/1183: member states must offer an EUDI Wallet to citizens, residents and businesses (deadline: December 2026); fewer than a third of states meet the readiness bar today (Politecnico di Milano observatory). https://ec.europa.eu/digital-building-blocks/sites/spaces/EUDIGITALIDENTITYWALLET/pages/915931811/The+European+Digital+Identity+Regulation · https://www.corbado.com/blog/eudi-wallet-2026-deadline-rollout-eic-2026
645. t2p — product manual (video-selfie signature: "the app's strongest signature" — the front camera records while you sign, a natural blink and two random liveness challenges (turn the head, smile, tilt it) in under 60 seconds, mandatory device biometric verification with PIN as an alternative, handwritten signature, 2 credits and two transactions; dedicated certificate with the document, the signature, the video evidence and the two transactions; only one person in the frame; certificate with SHA-512, RFC 3161/eIDAS seal, anchor to the public blockchain, QR and forensic check with device integrity and hardware attestation; ETSI .asice package verifiable offline without trusting the manufacturer; files stay on the device — only the fingerprint travels; the proof of life confirms "a real, present person" at signing). See chapters 13 and 15; executive impersonation and the "no seal, no payment" rule, in chapter 87; who you hire remotely, in chapter 91; commit-and-reveal, in chapters 45 and 46; the unique person of participation, in chapter 108.
646. 340B Report — *Fake Lilly Twitter account tweets 'Insulin Is Free Now,' temporarily wiping billions from company's stock value* (Nov. 10, 2022, 1:36 p.m. ET; fall from ~$368 to ~$346; ~$15 billion erased; the company's apology and the @LillyPad account). https://340breport.com/fake-lilly-twitter-account-tweets-insulin-is-free-now-temporarily-wiping-billions-from-companys-stock-value/
647. University at Buffalo — *With one false tweet, computer-based 'Hack Crash' led to real panic* (April 23, 2013; hacked AP tweet; Dow –143.5 points; >$136 billion erased in seconds; algorithmic and high-frequency reaction). https://www.buffalo.edu/news/releases/2015/05/037.html
648. InvestmentNews — *Shareholders accuse Ostin Technology of deepfake-driven $950M pump-and-dump scheme* (2025; deepfakes of Kostin/Musk/Zuckerberg; stolen identities of FINRA advisers; 1,175% inflation; 94% one-day collapse; DOJ indictment Sept. 2025). https://www.investmentnews.com/regulation-legal-compliance/legal-shareholders-accuse-ostin-technology-of-deepfake-driven-950m-pump-and-dump-scheme/265308
649. The Hans India — *BSE warns investors about fake deepfake video misusing CEO's identity* (2026 warnings over deepfake videos of chief executive Sundararaman Ramamurthy; recurring incidents January-April; promised legal action). https://www.thehansindia.com/business/bse-warns-investors-about-fake-deepfake-video-misusing-ceos-identity-1038779
650. SEC / FINRA / NASAA — *Investor Alert: Artificial Intelligence (AI) and Investment Fraud* (Investor.gov, Jan. 25, 2024; explicit warning about deepfake videos of chief executives to "manipulate a stock's price"). https://www.investor.gov/introduction-investing/general-resources/news-alerts/alerts-bulletins/investor-alerts/artificial-intelligence-fraud
651. PRNEWS — *Hoax Releases Hit Newswire Services* (fake Walmart-Litecoin release via GlobeNewswire; Litecoin +30%; the wire-authentication angle). https://www.prnewsonline.com/hoax-releases-hit-newswire-services/

652. Sidley Austin — *SEC Enforcement FY2025 Results Signal Shift in Priorities* (456 actions; $17.9 billion ordered; adjusted $1.4 billion + $1.3 billion; priority on market manipulation and disclosure fraud). https://www.sidley.com/en/insights/newsupdates/2026/04/sec-enforcement-fy2025-results-signal-shift-in-priorities-in-direct-critique-of-prior-administration
653. EUR-Lex / Lexparency — *MAR Regulation (EU) No 596/2014, Article 12 — Market manipulation* (dissemination of false or misleading information and rumors through the media and the internet as market manipulation). https://lexparency.org/eu/32014R0596/ART_12/
654. FINRA — *Artificial Intelligence (AI) and Investment Fraud* (FINRA's version of the alert; language on deepfake videos of chief executives and fake news to move a stock's price). https://www.finra.org/investors/insights/artificial-intelligence-and-investment-fraud
655. Carissa Véliz — *Prophecy: Prediction, Power, and the Fight for the Future* (2025).
656. ICIJ — *About the Pandora Papers investigation* (October 3, 2021; 11.9 million documents, 2.94 TB, 14 providers, 29,000+ beneficial owners, 35 leaders and 330+ politicians). https://www.icij.org/investigations/pandora-papers/about-pandora-papers-investigation/
657. Wikipedia — *Panama Papers* (April 3, 2016; 11.5 million documents from Mossack Fonseca, 214,488 entities, 150 investigations in 79 countries, shutdown in 2018). https://en.wikipedia.org/wiki/Panama_Papers
658. Transparency International — *Panama Papers: 10 years on* (~1,360M USD recovered; CTA reforms, EU registers, ROE; "the architecture of offshore secrecy remains intact"). https://www.transparency.org/en/news/panama-papers-10-years
659. FinCEN — *FinCEN Removes Beneficial Ownership Reporting Requirements for U.S. Companies* (interim rule of March 2025; exemption of domestic companies; covered population < 12,000). https://www.fincen.gov/news/news-releases/fincen-removes-beneficial-ownership-reporting-requirements-us-companies-and-us
660. Journal of Accountancy — *BOI update: 6.5 million of 32 million reports filed so far* (original estimate of 32 million obliged companies; 6.5M filed by November 2024). https://www.journalofaccountancy.com/news/2024/nov/boi-update-6-5-million-of-32-million-reports-filed-so-far/
661. Transparency International — *EU court ruling on beneficial ownership registers: legitimate access* (CJEU judgment of November 22, 2022; invalidation of public access; the legitimate interest framework). https://www.transparency.org/en/blog/eu-court-ruling-on-beneficial-ownership-registers-legitimate-access
662. PwC Ireland — *The EU's new Anti-Money Laundering Authority* (AMLA Frankfurt operational since July 1, 2025; AMLR (EU) 2024/1624 applicable July 10, 2027; interoperable UBO registers). https://www.pwc.ie/services/audit-assurance/insights/eu-new-anti-money-laundering-authority.html
663. GOV.UK — *Companies House confirms identity verification rollout from 18 November 2025* (ECCTA; identity verification of directors and PSCs; 100,400 companies affected, 10,200+ filings rejected). https://www.gov.uk/government/news/companies-house-confirms-identity-verification-rollout-from-18-november-2025
664. Transparency International US — *Unanimous Eleventh Circuit decision upholds the Corporate Transparency Act* (December 16, 2025; the highest court to uphold the CTA's constitutionality). https://us.transparency.org/news/unanimous-eleventh-circuit-decision-upholds-the-corporate-transparency-act/
665. FATF — *Beneficial Ownership* (Recommendation 24 revised in 2022 and 2023 guidance; adequate, accurate and up-to-date information; a multi-pronged approach). https://www.fatf-gafi.org/en/topics/beneficial-ownership.html
666. Open Ownership — *Verification of beneficial ownership data: Overview* (self-declared data needs verification: compliance, cross-checks, supporting evidence). https://www.openownership.org/en/publications/verification-of-beneficial-ownership-data/overview/
667. KnowBe4 — *How a North Korean Fake IT Worker Tried to Infiltrate Us* (Stu Sjouwerman; a first-person account: AI-retouched stock photo, four video interviews, EDR detection on July 15, 2024, mule laptop farm, FBI/Mandiant). https://blog.knowbe4.com/how-a-north-korean-fake-it-worker-tried-to-infiltrate-us
668. FBI/IC3 — *Public Service Announcement I-012325-PSA: North Korean IT Workers Conducting Data Extortion* (January 23, 2025; AI face-swapping in interviews, the pivot to extortion, identity verification recommendations). https://www.ic3.gov/PSA/2025/PSA250123

669. US DOJ (USAO-DC) — *Arizona Woman Sentenced in $17M IT Worker Fraud Scheme for North Korea* (July 24, 2025; Christina Chapman, 102 months, 309 companies, 68 stolen identities, $17 million, laptop farm, 49 laptops shipped abroad). https://www.justice.gov/usao-dc/pr/arizona-woman-sentenced-17m-it-worker-fraud-scheme-illegally-generated-revenue-north
670. Crowell & Moring — *DOJ Announces Major Enforcement Actions Targeting North Korean Remote IT Worker Schemes* (November 14, 2025; more than 136 victim companies, more than 18 compromised identities, more than $2.2 million in proceeds). https://www.crowell.com/en/insights/client-alerts/doj-announces-major-enforcement-actions-targeting-north-korean-remote-it-worker-schemes
671. The Decoder — *Gartner predicts that by 2028, one in four job applicant profiles will be fake* (the 25% forecast for 2028; the 6% admission of interview impersonation, survey of 3,000 candidates). https://the-decoder.com/gartner-predicts-that-by-2028-one-in-four-job-applicant-profiles-will-be-fake/
672. People Management — *Deepfakes and AI-enabled impersonation rank among top recruitment threats* (January 19, 2026; First Advantage 69% of UK leaders; Checkr 23% identity fraud; Greenhouse 91%). https://www.peoplemanagement.co.uk/article/1945557/deepfakes-ai-enabled-impersonation-rank-among-top-recruitment-threats-research-reveals
673. Pindrop — *The Growing Trend of Deepfakes in Interviews* (AI audio in ~1 in 106 meetings; ~1,300% year-on-year; synthetic voices in 0.33% of calls in Q4 2024). https://www.pindrop.com/article/growing-trend-of-deepfakes-in-interviews/
674. NIST — *SP 800-63-4 Digital Identity Guidelines* (final, July 2025; IAL1/2/3 assurance levels and the Remote/In-Person × Attended/Unattended verification taxonomy). https://pages.nist.gov/800-63-4/sp800-63.html
675. FBI — *Fraudulent Remote IT Workers from DPRK* (wanted/guidance; use of stolen identities, intermediary computers and AI, with mitigations for employers). https://www.fbi.gov/wanted/cyber/fraudulent-remote-it-workers-from-dprk
676. Microsoft Security — *Fake Employees, Real Threat: Decentralized Identity to combat Deepfake Hiring* (verifiable credentials and continuous verification of onboarding identity; adjacent-solution context). https://techcommunity.microsoft.com/blog/microsoft-security-blog/fake-employees-real-threat-decentralized-identity-to-combat-deepfake-hiring/4479864
677. New York Attorney General (2021). https://ag.ny.gov/press-release/2021/attorney-general-james-issues-report-detailing-millions-fake-comments-revealing
678. OJEU / BOE — *Directive (EU) 2024/2831 of the European Parliament and of the Council, of October 23, 2024, on improving working conditions in platform work* (in force December 2, 2024; transposition by December 2, 2026; presumption of employment on facts of direction and control with the burden of proof on the platform; ban on sensitive data in automated decisions; transparency, human oversight and human intervention in important decisions); Laboral Social — *Directiva sobre plataformas digitales de trabajo*. https://www.boe.es/buscar/doc.php?id=DOUE-L-2024-81667 · https://eur-lex.europa.eu/eli/dir/2024/2831/oj?eliuri=eli:dir:2024:2831:oj&locale=es · https://www.laboral-social.com/directiva-plataformas-digitales-trabajo-mejora-condiciones-laborales
679. BOE — *Law 12/2021, of September 28* (the "Rider Law": presumption of employment in delivery via algorithmic management; the works council's right to be informed of the parameters, rules and instructions of the algorithms affecting working conditions; in force since August 2021); PwC Periscopio — *¿Ha cumplido sus objetivos la Ley Rider?*. https://www.boe.es/buscar/act.php?id=BOE-A-2021-15767 · https://periscopiofiscalylegal.pwc.es/luces-y-sombras-de-la-ley-rider/
680. elDiario.es — *El jefe de Glovo admite a la jueza que la regularización de los 'riders' busca evitar «discusiones» con Inspección de Trabajo* (criminal proceedings in Barcelona; ≈€200 million in contributions and fines; Just Eat claims €295 million); elDiario.es — *El jefe de Glovo dice que la presión penal y las multas forzaron cambiar su modelo laboral: «La caja no te da para operar»* (10,000-15,000 riders; Delivery Hero provisions of €440-770 million; ≈€100 million impact on 2025 EBITDA); Crónica Global — *Óscar Pierre justifica el cambio de modelo laboral por «presión económica y penal»*. https://www.eldiario.es/catalunya/jefe-glovo-admite-jueza-regularizacion-riders-busca-evitar-discusiones-inspeccion-trabajo_1_11870197.html · https://www.eldiario.es/economia/jefe-glovo-dice-presion-penal-multas-forzaron-cambiar-modelo-laboral-caja-no-da-operar_1_12062602.html · https://cronicaglobal.elespanol.com/business/20250218/oscar-pierre-de-glovo-justifica-laboral-economica/925157487_0.html

681. Fountain Court Chambers — *Amsterdam Court Upholds Appeal in Algorithmic Decision-Making Test Case: Drivers v Uber and Ola* (April 2023; three judgments; Ola must explain the logic of the fully automated decision, Art. 22 GDPR); Worker Info Exchange — *Uber ordered to pay €584,000 for failure to comply with court order in robo-firing case* (plus €4,000/day); UK Human Rights Blog — *Amsterdam Court orders reinstatement of Uber drivers dismissed by algorithm*. https://fountaincourt.uk/2023/04/amsterdam-court-upholds-appeal-in-algorithmic-decision-making-test-case-drivers-v-uber-and-ola/ · https://www.workerinfoexchange.org/post/uber-ordered-to-pay-584-000-for-failure-to-comply-with-court-order-in-robo-firing-case/ · https://ukhumanrightsblog.com/2021/05/18/amsterdam-court-orders-reinstatement-of-uber-drivers-dismissed-by-algorithm/
682. t2p — product manual (SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; content on the device, voluntary disclosure; the series' gaps get dated). See chapters 13 and 15; the provenance of AI models, in chapter 79.
683. Ministry of Labor / trade press — *Registro horario digital obligatorio: la reforma en trámite* (urgent processing approved by the Council of Ministers on 30-9-2025, decoupled from the working-week reduction; aim of approval during 2026, no entry-into-force date published; the draft's requirements: traceable, immutable electronic records, timestamping, recording of breaks, remote real-time access for the Labor Inspectorate, a history of modifications with author and date); EPA/INE — more than 40% of overtime goes unpaid; RD-ley 8/2019 — mandatory daily working-time record. https://www.protime.eu/es-es/noticias/registro-horario-digital-obligatorio-2026 · https://factorial.es/blog/nuevo-registro-horario-digital-espana/ · https://www.inwout.com/post/registro-horario-digital-obligatorio-2026 · https://kronjop.com/es/newsroom/control-horario/digital/obligatorio/
684. The general labor framework (the labor court claim; the Labor Inspectorate's reports; the delegates' role). See chapter 92 for platform work and chapter 102 for the labor supply chain.
685. t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; certified screenshot and email; certify any file — only the fingerprint travels; background certification; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; one's own conversations and harassment, in chapter 75.
686. The Dial — *Faulty Construction* (the Turkey earthquake trials; the Isias Grand Hotel; 18-year sentences in December 2024; 75,000 amnestied buildings; some 13 million noncompliant buildings). https://www.thedial.world/articles/news/issue-25/turkey-earthquake-trials
687. PBS Peril & Promise — *How unenforced building codes worsened the Turkey-Syria earthquake* (July 2023; amnesties, unenforced codes and concrete quality failures). https://www.pbs.org/wnet/peril-and-promise/2023/07/unenforced-building-codes-worsened-the-impacts-of-the-earthquake-that-hit-turkey-and-syria/
688. NIST — *Champlain Towers South Collapse* (the official federal investigation; 98 dead; initiation by punching shear failure of the pool deck slab; concrete weaker than the code). https://www.nist.gov/disaster-failure-studies/champlain-towers-south-collapse
689. ICAC Hong Kong — *Consultant director jailed over false HZMB concrete cube test records* (19 workers charged; 384 falsified reports; date manipulation and cube substitution; Albert Leung, 2.5 years in June 2024). https://www.icac.org.hk/en/p/press/index_id_1890.html
690. The Real Deal — *Concrete tester exec convicted of falsifying safety reports for Yankee Stadium, Freedom Tower* (February 2010; Testwell/Kancharla; the Freedom Tower certified at 12,000 PSI but around 9,000 PSI; 7 to 21 years). https://therealdeal.com/new-york/2010/02/17/concrete-tester-exec-convicted-of-falsifying-safety-reports-for-yankee-stadium-freedom-tower/
691. Wikipedia — *Ponte Morandi collapse* (Genoa, August 14, 2018; 43 dead; 59 defendants; "clearly false" inspection documents; self-surveillance via Spea; the €30 million settlement in 2023). https://en.wikipedia.org/wiki/Ponte_Morandi_collapse
692. Seton UK — *Understanding the Golden Thread: Building Safety Act 2022* (a traceable, version-controlled digital record for higher-risk buildings, post-Grenfell). https://www.seton.co.uk/legislationwatch/article/understanding-the-golden-thread-the-building-safety-act-2022/
693. EUR-Lex — *Regulation (EU) 2024/3110* (the revised Construction Products Regulation; in force January 7, 2025; the Digital Product Passport mandatory for CE marking from January 8, 2026). https://eur-lex.europa.eu/eli/reg/2024/3110/oj/eng

694. NIST — *Report to Congress provides update on Champlain Towers South investigation* (February 2025; experimental work complete, modeling underway, technical presentation in June 2025). https://www.nist.gov/news-events/news/2025/02/nist-report-congress-provides-update-champlain-towers-south-investigation
695. BSI — *ISO 19650 Building Information Modelling standards* (information management with BIM and the as-built record / digital twin across the life cycle). https://www.bsigroup.com/en-GB/products-and-services/standards/iso-19650-building-information-modelling-bim/
696. SFO (GOV.UK) — *SFO secures 4-year prison sentence for aircraft parts fraud* (Zamora Yrala; 60,000 parts; a £39.3 million fraud; forged ARCs; a 4-year-8-month sentence; discovery by Safran). https://www.gov.uk/government/news/sfo-secures-4-year-prison-sentence-for-aircraft-parts-fraud
697. SFO (GOV.UK) — *SFO case page: AOG Technics Ltd* (the Serious Fraud Office's official case page). https://www.gov.uk/sfo-cases/aog-technics-ltd
698. Fortune — *Fake components went into 68 jet engines on Boeing 737 and Airbus A320 planes, lawsuit says* (CFM International's suit: 78 falsified documents, 52 CFM56 part numbers, 68 engines, turbine blades). https://fortune.com/2023/09/08/fake-components-68-jet-engines-boeing-737-airbus-a320-planes-lawsuit/
699. FAA — *Suspected Unapproved Parts (SUP) Program* (traceability under 14 CFR Part 21/43, Form 8130-3, reporting with Form 8120-11). https://www.faa.gov/aircraft/safety/programs/sups
700. EASA — *Suspected Unapproved Parts* (the EU database and definitions; cases of forged EASA Form 1; the SIB framework). https://www.easa.europa.eu/en/domains/aircraft-products/suspected-unapproved-parts
701. EASA — *Safety Information Bulletin SIB 2023-05* (the bulletin on AOG Technics's suspect parts and the recommended quarantine/replacement actions). https://ad.easa.europa.eu/blob/EASA_SIB_2023_05.pdf/SIB_2023-05_1
702. GE Aerospace — *Coalition delivers report to help prevent future unapproved parts entering aviation* (the Aviation Supply Chain Integrity Coalition: 13 recommendations, October 9, 2024; digitize 8130-3/Form 1; traceability; supplier accreditation). https://www.geaerospace.com/news/press-releases/coalition-delivers-report-help-prevent-future-unapproved-parts-entering-aviation
703. Connector Supplier — *Counterfeit Components Ground Airlines* (~520,000 counterfeit or unapproved parts reach aircraft each year, ~2% of 26 million; a market above $2 billion). https://connectorsupplier.com/counterfeit-components-ground-airlines/
704. PR Newswire / CARFAX — *Odometer rollbacks see dramatic jump of 14% to 2.45 million vehicles* (December 2025; 2.45 million suspect vehicles; +14% year-on-year; ~$3,300 average loss). https://www.prnewswire.com/news-releases/carfax-odometer-rollbacks-see-dramatic-jump-of-14-to-2-45-million-vehicles-302643855.html
705. EPA — *Volkswagen Clean Air Act Civil Settlement* (2016-2017 settlements; a $1.45 billion civil penalty; the mitigation fund; ~590,000 U.S. vehicles; the defeat device description). https://www.epa.gov/enforcement/volkswagen-clean-air-act-civil-settlement
706. Wikipedia — *Volkswagen emissions scandal* (NOx up to 40 times the limit; ~11 million cars; ~$33 billion total cost; the global chronology). https://en.wikipedia.org/wiki/Volkswagen_emissions_scandal
707. EPA — *United States and California Announce Diesel Engine Manufacturer Cummins Inc. Agrees to Pay Record $1.675 Billion Civil Penalty* (January 10, 2024; the highest civil penalty under the Clean Air Act; 630,000+ RAM trucks; $325 million in remedies; the recall conditions). https://www.epa.gov/newsreleases/united-states-and-california-announce-diesel-engine-manufacturer-cummins-inc-agrees
708. EPA — *Summary of Criminal Prosecutions: Berkshire Power* (intentional CEMS biasing, 2009-2011; the fraudulent adjustment removed before the audits; a $2.75 million criminal fine for BPC; ~$7.25 million total by March 2017). https://cfpub.epa.gov/compliance/criminal_prosecution/index.cfm?action=3&prosecution_summary_id=2981
709. IEA — *Global Methane Tracker 2024, Key Findings* (close to 120 Mt of fossil methane in 2023; UNFCCC ~50% lower; OGMP 2.0 ~95% lower; satellite-detected leak events +50% year-on-year). https://www.iea.org/reports/global-methane-tracker-2024/key-findings
710. MethaneSAT (EDF) — *New Data Show US Oil and Gas Methane Emissions Over Four Times Higher Than EPA Estimates* (measured emissions >4 times the EPA's estimates; a 1.6% loss rate, eight times the 0.2% target; 30+ flights covering ~70% of U.S. onshore production).

https://www.methanesat.org/project-updates/new-data-show-us-oil-and-gas-methane-emissions-over-four-times-higher-epa-estimates

711. EPA — *EPA and Justice Department Announce $241.5M Settlement with Marathon Oil* (July 11, 2024; a $64.5 million civil penalty, the highest for a stationary source; ~90 facilities; ~110,000 tonnes of VOCs/methane cut). https://www.epa.gov/newsreleases/epa-and-justice-department-announce-2415m-settlement-marathon-oil-reduce-climate-and
712. CJEU — *Press release on the July 2022 judgment that thermal-window defeat devices are in principle illegal* (PDF; the legal definition of the defeat device; the earlier December 2020 judgment on test-cycle recognition). https://curia.europa.eu/site/upload/docs/application/pdf/2022-07/cp220124en.pdf
713. EPA — *Key Things to Know About EPA's Final Methane Rule* (fact sheet PDF; 58 million tonnes of methane avoided 2024-2038; a ~80% reduction; the Super-Emitter Program). https://www.epa.gov/system/files/documents/2023-12/key-things-to-know-about-epas-final-rule-for-oil-and-natural-gas-operations.fact-sheet.pdf
714. European Commission — *New EU Methane Regulation to reduce emissions from fossil fuels* (Regulation (EU) 2024/1787; source-level MRV with third-party verification; satellite super-emitter rapid alert; MRV equivalence for importers from 2027). https://energy.ec.europa.eu/news/new-eu-methane-regulation-reduce-harmful-emissions-fossil-fuels-europe-and-abroad-2024-05-27_en
715. El Ágora Diario — *Europa condena a España por no proteger el agua y hábitats de Doñana* (the CJEU judgment of June 24, 2021; illegal extractions not counted in the 2015-2021 Hydrological Plan; complaints since 2009); Greenpeace — *Doñana, agua sin ley* (wells, ponds and pipes; strawberry greenhouses; deforestation). https://www.elagoradiario.com/agua/tribunal-justicia-europa-condena-espana-agua-donana/ · https://es.greenpeace.org/es/en-profundidad/sos-acuiferos/donana-agua-sin-ley/
716. BOE — *Law 19/2022, of September 30, recognizing legal personhood for the Mar Menor lagoon and its basin* (the first ecosystem in Europe; rights to exist, evolve, protection, conservation and restoration; a popular initiative with 639,826 signatures); Araoz & Rueda — *El TC reconoce la personalidad jurídica del Mar Menor*; BOE — *Royal Decree 90/2025* (committee of representatives: 13 members, 7 from civil society); Newtral — *Qué significa que el Mar Menor haya adquirido personalidad jurídica*. https://www.boe.es/buscar/doc.php?id=BOE-A-2022-16019 · https://www.araozyrueda.com/el-tribunal-constitucional-reconoce-mar-menor-sujeto-de-derechos/ · https://www.boe.es/diario_boe/txt.php?id=BOE-A-2025-2577 · https://www.newtral.es/mar-menor-personalidad-juridica/20220923/
717. Sensor.Community / World Air Quality Index — *Real-time citizen air quality sensor network (formerly luftdaten.info)* (more than 100 countries; low-cost citizen sensors); AireCiudadano — an open network of particle monitors with a map and API; SciStarter/CanAirIO — open PM2.5 sensors. https://aqicn.org/network/luftdaten/es/ · https://aireciudadano.com/ · https://scistarter.org/crea-tu-propio-sensor-de-calidad-del-aire-con-cana
718. Perito Judicial — *Ruidos molestos y su reclamación: la clave está en la medición* (municipal inspection with a homologated, calibrated sound level meter; reports with a presumption of veracity and full evidentiary value; indicative indoor limits 35-45 dB day / 25-30 dB night); COITTA — *La importancia de la prueba pericial acústica*; the Ombudsman — *Comprobaciones ante las denuncias de ruidos molestos* (providing the means to check complaints). https://peritojudicial.com/ruidos-molestos/ · https://www.aagit.org/la-importancia-de-la-prueba-pericial-acustica-y-la-proteccion-juridica-frente-al-ruido-como-agente-contaminante/ · https://www.defensordelpueblo.es/resoluciones/habilitar-los-medios-para-poder-realizar-comprobaciones-necesarias-cuando-se-denuncien-ruidos-molestos/
719. t2p — product manual (SHA-512 fingerprint at the origin — sensor or device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; .asice container; real-time batch series; dated gaps; content on the device and voluntary disclosure). See chapters 13 and 15; the industrial sensor, in chapter 96; the orbital witness, in chapter 104.
720. EUR-Lex — *Sustainability rules for batteries and waste batteries (Regulation (EU) 2023/1542 summary)* (the passport from 2027 via QR and unique identifier; recycled content targets from August 18, 2031: Co 16%, Pb 85%, Li 6%, Ni 6%; customs access to the registry). https://eur-lex.europa.eu/EN/legal-content/summary/sustainability-rules-for-batteries-and-waste-batteries.html
721. DigiProdPass — *EU Battery Regulation: Recycled Content Targets (2031-2036)* (the 2031 percentages — Co 16/Pb 85/Li 6/Ni 6 — and the 2036 ones — Co 26/Pb 85/Li 12/Ni 15 — the

falsifiable claims the passport carries). https://digiprodpass.com/blogs/eu-battery-regulation-recycled-content-targets

722. Minespider — *EU Battery Due Diligence Obligations postponed* (due diligence postponed from Aug 18, 2025 to Aug 18, 2027 via the "stop-the-clock" Regulation (EU) 2025/1561; guidance extended until Jul 26, 2026). https://www.minespider.com/blog/eu-battery-due-diligence-postponed-what-it-means-for-companies
723. Electrek — *Volvo EX90 launches world's first EV battery passport* (Jun 4, 2024; five years of development with Circulor; ~$10/vehicle for 15 years; the QR on the door pillar; traces materials, recycled content and carbon footprint). https://electrek.co/2024/06/04/volvo-ex90-launch-worlds-first-ev-battery-passport/
724. IEEE Spectrum — *Volvo's Battery Passport Traces the Supply Chain* (more than 20 billion data points across 150 million cells; ~52% of the world's cell makers; 15 years of state-of-health monitoring). https://spectrum.ieee.org/volvo-ex90-battery-passport
725. Global Battery Alliance — *World's First Battery Passport Proof of Concept* (Davos, Jan 18, 2023; Audi/BASF/CATL/Glencore/Tesla/VW; the data "has not been fully verified and was partially complemented with reasonable estimates"). https://www.globalbattery.org/press-releases/global-battery-alliance-launches-world%E2%80%99s-first-battery-passport-proof-of-concept/
726. European Commission (Green Forum) — *Implementing the Ecodesign for Sustainable Products Regulation (ESPR)* (the ESPR in force Jul 18, 2024; the digital product passport as a horizontal tool; the 2025-2030 work plan adopted Apr 16, 2025). https://green-forum.ec.europa.eu/implementing-ecodesign-sustainable-products-regulation_en
727. Hogan Lovells — *Digital Product Passports in the EU: expansion under ESPR and lessons from the Battery Passport pilot* (legal analysis distinguishing the battery pilot from the ESPR's horizontal rollout; the delegated acts calendar). https://www.hoganlovells.com/en/publications/digital-product-passports-in-the-eu-comprehensive-expansion
728. CIRPASS — *Standardisation gaps and roadmap (v1.2, July 2024)* (the DPP data model, unique identifiers, data carriers and the JTC 24 standardization gaps — the architecture the seal plugs into). https://cirpassproject.eu/wp-content/uploads/2024/07/CIRPASS-standardisation-gaps-and-roadmap-V1.2.pdf
729. OJEU/BOE — *Regulation (EU) 2024/1157, of April 11, 2024, on shipments of waste* (repeals Regulation 1013/2006; generally applicable from May 21, 2026; exports of all plastic waste to third countries subject to prior written notification and consent; a ban on exporting waste for recovery to non-OECD countries; electronic processing); RETEMA — *Nuevo Reglamento de traslados de residuos*; Residuos Profesional — *Análisis del Reglamento (UE) 2024/1157* (context: Turkey received 39% of European waste exports — 12.4 million tonnes, 2022 data). https://www.boe.es/buscar/doc.php?id=DOUE-L-2024-80606 · https://eur-lex.europa.eu/legal-content/es/ALL/?uri=CELEX:32024R1157 · https://www.retema.es/blog/nuevo-reglamento-de-traslados-de-residuos-la-normativa-que-pone-en-jaque-la-industria-europea · https://www.residuosprofesional.com/reglamento-relativo-traslados-de-residuos/
730. The general waste framework (identification and tracking documents for shipments; extended producer responsibility; the waste hierarchy). See chapter 98 for the digital product passport and chapter 96 for emissions.
731. t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; video with date/time overlay; document scanner up to 50 pages; certify any file in batches — only the fingerprint travels; web page capture with full image + source code; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the product passport, in chapter 98; citizen science, in chapter 97.
732. Agrodigital — *El cuaderno digital de explotación será obligatorio a partir del 1 de enero de 2028* (the CUE voluntary since 1-1-2024; postponements; mandatory from 1-1-2028 for holdings >30 ha of crops, >5 ha irrigated or >0.1 ha greenhouse; the electronic plant protection record mandatory from 1-1-2027 — paper until 31-12-2026); FEGA/Agronews — entry-into-force clarifications; Government of Aragon — *Cuaderno digital (CUE)* (content: plant protection operations, fertilization, irrigation and other practices). https://agrodigital.info/2025/01/23/el-cuaderno-digital-de-explotacion-sera-obligatorio-a-partir-del-1-de-enero-de-2028/ · https://www.agronewscastillayleon.com/el-fega-publica-aclara-la-entrada-en-vigor-del-cuaderno-de-explotacion-y-el-plan-de-abonado/ · https://www.aragon.es/pac-2026/cuaderno-digital-cue · https://www.fega.gob.es/sites/default/files/files/document/10.2023.03.15-12.45h_blq-04_inmaculada-bravo_cue.el-cue-y-el-riego.pdf

733. CAP 2023-2027 — the monitoring system (Sentinel/Copernicus imagery over the parcels declared in the SIGPAC; requests for georeferenced photographs from the farmer when the system cannot conclude). The FEGA's and the European Commission's general framework. See chapter 104 for the integrity of the orbital witness.
734. t2p — product manual (photo with forensic GPS — genuine fix ~50 m / ≤30 s, a simulated location blocks certification; background certification — the field without coverage; certify any file in batches; document scanner; projects for a season's file; SHA-512 fingerprint on the device; anchoring of the fingerprint to the public blockchain; complementary RFC 3161/eIDAS timestamp; PDF certificate with QR; .asice package verifiable offline). See chapters 13 and 15; the satellite, in chapter 104; food, in chapter 101; waste, in chapter 99.
735. Europol — *EUR 91 million worth of counterfeit and substandard food seized in Europe-wide operation* (Operation OPSON XIII; 22,000 t of food, €91 million, 11 networks; Italy: extra virgin and chlorophyll; Spain: oil and tuna). https://www.europol.europa.eu/media-press/newsroom/news/eur-91-million-worth-of-counterfeit-and-substandard-food-seized-in-europe-wide-operation
736. Europol — *Counterfeit and substandard food worth EUR 95 million seized in global operation* (Operation OPSON XIV; €95 million, 11,566 t, 631 reported, 13 organized crime groups, 31 countries). https://www.europol.europa.eu/media-press/newsroom/news/counterfeit-and-substandard-food-worth-eur-95-million-seized-in-global-operation
737. CNN — *Italian authorities confiscate almost $1 million in fake olive oil* (July 11, 2024; the Carabinieri raid in Apulia, 42 t of fake extra virgin, 623 L of chlorophyll, 7 charged, forged labels and stamps). https://www.cnn.com/2024/07/11/europe/italy-olive-oil-seizure-scli-intl/
738. Food Safety News — *Spanish and Italian investigators uncover olive oil fraud* (December 2023; Operation Omegabad: 260,000 L seized, 11 arrests, lampante sold as extra virgin, Ciudad Real/Sicily/Tuscany, €91,000 in cash). https://www.foodsafetynews.com/2023/12/spanish-and-italian-investigators-uncover-olive-oil-fraud/
739. European Commission (DG SANTE) — *EU coordinated action on honey ('From the Hives')* (2021-2022 samples; 46% of 320 consignments suspected of adulteration with sugar syrup; methodology and participating countries). https://food.ec.europa.eu/food-safety/eu-agri-food-fraud-network/eu-coordinated-actions/honey-2021-2022_en
740. FoodNavigator — *Geographical Indications worth a mouth-watering €75 billion, study finds* (the value of EU GIs: €74.8 billion in 2017; a price premium of about 2× for protected products). https://www.foodnavigator.com/Article/2020/04/22/Geographical-Indications-worth-a-mouth-watering-75-billion-study/
741. European Commission — *EU Agri-Food Fraud Network: reports and publications* (fraud affects ~1% of the industry and costs $10,000-40,000 million/year; 758 suspicions in 2023, +26%; meat, honey and oil at the top). https://food.ec.europa.eu/food-safety/eu-agri-food-fraud-network/reports-and-publications_en
742. EUR-Lex — *Regulation (EU) 2024/1143 on geographical indications* (in force May 13, 2024; a single PDO/PGI/GI and TSG framework; protection as an ingredient and online). https://eur-lex.europa.eu/eli/reg/2024/1143/oj/eng
743. Food Safety News — *Four sentenced in France for roles in horse meat scandal* (the 2019 trial; Poujol (Spanghero) 2 years + a €100,000 fine, Fasen 2 years + a permanent ban; Comigel, 500+ t to 13 countries). https://www.foodsafetynews.com/2019/04/four-sentenced-in-france-for-roles-in-horse-meat-scandal/
744. FDA — *Economically Motivated Adulteration (Food Fraud)* (the EMA definition, e.g. cheap oil sold as 100% olive oil; the FSMA's preventive controls duty). https://www.fda.gov/food/compliance-enforcement-food/economically-motivated-adulteration-food-fraud
745. Sensors (MDPI) — *Blockchain and IoT for food traceability in olive oil supply chains* (2024; the state of the art in sealing/anchoring oil provenance and in NMR/IRMS analytical authentication). https://www.mdpi.com/1424-8220/24/24/8189
746. ILO — *Data and research on forced labour (Global Estimates 2022)* (27.6 million people in forced labor in 2021; 17.3 million in the private economy; migrants at triple the risk). https://www.ilo.org/topics/forced-labour-modern-slavery-and-trafficking-persons/data-and-research-forced-labour
747. ILO — *Annual profits from forced labour amount to US$236 billion* (March 2024; +37% since 2014; $84 billion in Europe and Central Asia). https://www.ilo.org/resource/news/annual-profits-forced-labour-amount-us-236-billion-ilo-report-finds

748. Miller & Chevalier — *UFLPA Enforcement 2024 Year in Review* (16,755 shipments and $3.69 billion detained; $1.34 billion in 2024; 48% denied entry; +1,600% in automotive; the *Entity List* at 144). https://www.millerchevalier.com/publication/trade-compliance-flash-uflpa-enforcement-2024-year-review
749. Business & Human Rights Resource Centre — *Volkswagen cars held up at US ports over part linked to allegations of forced labour in Xinjiang* (the brands, the sub-component, the December 2023 – mid-January 2024 chronology). https://www.business-humanrights.org/en/latest-news/volkswagen-cars-held-up-at-us-ports-over-part-linked-to-allegations-of-forced-labour-in-xinjiang/
750. Human Rights Watch — *Asleep at the Wheel: Car Companies' Complicity in Forced Labor in China* (February 2024; VW, BMW, Mercedes, GM and the aluminum chain). https://www.hrw.org/report/2024/02/01/asleep-wheel/car-companies-complicity-forced-labor-china
751. European Commission — *Corporate Sustainability Due Diligence (CSDDD)* (Directive 2024/1760, in force July 25, 2024; transposition by July 26, 2028, application from July 26, 2029 after the Omnibus). https://commission.europa.eu/business-economy-euro/doing-business-eu/sustainability-due-diligence-responsible-business/corporate-sustainability-due-diligence_en
752. Covington — *EU Forced Labour Regulation Enters Into Force* (December 2024; in force December 13, 2024, application December 14, 2027; an investigation model without a geographic presumption). https://www.cov.com/en/news-and-insights/insights/2024/12/eu-forced-labour-regulation-enters-into-force
753. BAFA — *Supply Chain Act (LkSG): Overview* (application to companies of 1,000 employees since January 1, 2024; fines up to 2% of worldwide turnover; 486 audits in 2023). https://www.bafa.de/EN/Supply_Chain_Act/Overview/overview_node.html
754. ECCHR — *More for show than safety: Certificates in the textile industry* (Rana Plaza, 1,131 dead, BSCI; Ali Enterprises, more than 250 dead, SA8000; both had passed inspection). https://www.ecchr.eu/en/case/more-for-show-than-safety-certificates-in-the-textile-industry/
755. Human Rights Watch — *Obsessed with Audit Tools, Missing the Goal: Why Social Audits Can't Fix Labor Rights Abuses* (November 2022; coached workers, double books, fake payrolls, hidden children). https://www.hrw.org/report/2022/11/15/obsessed-audit-tools-missing-goal/why-social-audits-cant-fix-labor-rights-abuses
756. COSH — *Living wages for EU garment workers* (the legal minimum wage is 31% of a living wage in India, 37% in China and Bangladesh, 33% in Cambodia, 13% in Sri Lanka). https://cosh.eco/en/articles/living-wages-eu-garment-workers
757. GPS spoofing is scrambling ships in the Strait of Hormuz — *Scientific American* (spoofing in Hormuz during the 2025 Israel-Iran war; ships on land and inside nuclear plants; Humphreys's attribution to Iran; ~20% of the world's oil through the strait). https://www.scientificamerican.com/article/gps-spoofing-is-scrambling-ships-in-the-strait-of-hormuz/
758. Iran's $800M oil-smuggling scheme uses tankers posing as Iraqi ships — *Fox News* (Windward's disclosure, April 29, 2026; VLCCs with IMO numbers; fake calls at Khor al Zubair). https://www.foxnews.com/world/irans-800m-oil-smuggling-scheme-uses-tankers-posing-iraqi-ships-dodge-blockade-.amp
759. AIS spoofing: the fast track to sanctions — *Kpler* (261 vessels faked their AIS before being sanctioned, January 2024-July 2025; fake calls at Khor al Zubair). https://www.kpler.com/blog/ais-spoofing-fast-track-to-sanctions
760. Maritime GNSS Interference Worldwide: A Cumulative Analysis 2025 — *GPSPATRON* (10,000+ vessels affected in Q2 2025, eight times more than Q1; regional hotspots). https://gpspatron.com/maritime-gnss-interference-worldwide-a-cumulative-analysis-2025/
761. Widespread GPS jamming hits 1,000-plus ships in the Middle East Gulf — *Windward* (March 1, 2026: 1,100+ vessels in 24 h, 21+ hotspots; false fixes at Barakah and Al Hamra; Hormuz down a third). https://windward.ai/blog/gps-jamming-disrupts-1100-ships-in-the-middle-east-gulf/
762. Russian Shadow Fleet Tracker — *Kyiv School of Economics (KSE) Institute* (the fleet's size; ~17% of the world's oil fleet; $9.4 billion in extra revenue in 2024; the growing weight of the Russian flag). https://kse.ua/about-the-school/news/russian-shadow-fleet-tracker-march-2026-share-of-russian-flagged-shadow-tankers-jumps-from-3-to-21-in-nine-months/
763. Council sanctions 41 vessels of the Russian shadow fleet — *Council of the EU (Consilium)* (December 18, 2025: the EU total reaches 444+ vessels; port access and service bans).

https://www.consilium.europa.eu/en/press/press-releases/2025/12/18/russia-s-war-of-aggression-against-ukraine-council-sanctions-41-vessels-of-the-russian-shadow-fleet/
764. EU adopts 18th package of economic and individual measures — *Council of the EU (Consilium)* (July 18, 2025: the crude cap lowered from $60 to $47.60 with a dynamic mechanism 15% below the market). https://www.consilium.europa.eu/en/press/press-releases/2025/07/18/russia-s-war-of-aggression-against-ukraine-eu-adopts-18th-package-of-economic-and-individual-measures/
765. Treasury Maintains Pressure on Iranian Shadow Fleet (MS ENOLA) — *U.S. Department of the Treasury / Iran Primer* (December 2024; OFAC action under E.O. 13902 citing the MS ENOLA "turning off the AIS"). https://iranprimer.usip.org/blog/2024/dec/19/us-sanctions-more-irans-dark-fleet
766. Estlink 2 incident / Eagle S shadow-fleet tanker — *NPR* (the Cook Islands-flagged tanker Eagle S, December 25, 2024, damage to the Baltic submarine cable; the hybrid threat dimension). https://www.npr.org/2024/12/31/nx-s1-5243302/finland-russia-severed-undersea-cable-shadow-fleet
767. AIS spoofing: vessel-identity manipulation typologies — *Lloyd's List Intelligence* (identity laundering, IMO number fraud, the AIS handshake and the mechanics of position spoofing). https://www.lloydslistintelligence.com/resources/blog/the-secret-lives-of-the-shadow-fleet-ais-spoofing-vessel-identity-manipulation
768. The Moscow Times — *Satellite Images Show Bodies in Bucha for Weeks, Rebutting Moscow Claim* (April 5, 2022; nine Maxar images of March 18/19/31; bodies on Yablonska Street; a 14-meter mass grave). https://www.themoscowtimes.com/2022/04/05/satellite-images-show-bodies-in-bucha-for-weeks-rebutting-moscow-claim-a77211
769. Aerospace America / AIAA — *Maxar describes its role in Bucha revelations* (the WorldView-3 capture; Daniel Jablonsky on Maxar's intensive collection; the New York Times's use on April 4, 2022). https://aerospaceamerica.aiaa.org/maxar-bucha-insights/
770. BuzzFeed News — *Blanked-Out Spots On China's Maps Helped Us Uncover Xinjiang's Camps* (the Pulitzer-winning method; Planet imagery and censored tiles; 268 new camps and 428 suspect sites). https://www.buzzfeednews.com/article/alison_killing/satellite-images-investigation-xinjiang-detention-camps
771. *Deepfake Geography: Detecting AI-Generated Satellite Images* (arXiv 2511.17766, November 2025; ViT 95.11% versus CNN 87.02% on more than 130,000 images; GAN versus diffusion artifacts). https://arxiv.org/abs/2511.17766
772. *Generating Synthetic Satellite Imagery With Deep-Learning Text-to-Image Models* (arXiv 2404.07754, April 2024; Nguyen, Glaser and Biessmann; photorealistic synthetic imagery of facilities including nuclear plants; arms control). https://arxiv.org/abs/2404.07754
773. European Commission — *Observer: How Galileo OSNMA helps counter GNSS spoofing* (OSNMA operational July 24, 2025; the world's first Open Service signal authentication of a GNSS via TESLA cryptographic signatures). https://defence-industry-space.ec.europa.eu/observer-how-galileo-osnma-helps-counter-gnss-spoofing-2025-09-08_en
774. SpaceNews — *NGA to launch $700 million program to help AI make sense of satellite images* (the Sequoia labeling contract of up to $708 million over 7 years, 2024; feeds the Maven computer vision; the agency's largest data contract). https://spacenews.com/nga-to-launch-700-million-program-to-help-ai-make-sense-of-satellite-images/
775. Open Geospatial Consortium — *Navigating Synthetic Imagery Trust in Geospatial Data* (the Integrity, Provenance and Trust initiative, April 2025; defines integrity/provenance/trust, but with no cryptographic signing or blockchain). https://www.ogc.org/blog-article/navigating-synthetic-imagery-trust-geospatial-data/
776. Maldita.es — *Quién envió la alerta de las 20:11 horas en Valencia, qué dice el protocolo* (October 31, 2024; ES-Alert at 20:11 and 21:03; at level 2 the request belongs to the regional government). https://maldita.es/malditateexplica/20241031/generalitat-valenciana-responsable-alerta-movil/
777. elDiario.es — *El envío de la alerta masiva a móviles se planteó en la reunión clave de la DANA pasadas las 18.00, dos horas antes* (the CECOPI convened at 15:00, met at 17:00; the mass alert proposed past 18:00). https://www.eldiario.es/comunitat-valenciana/reunion-clave-dana-debatio-enviar-alerta-masiva-moviles-pasadas-18-00-horas_1_11802363.html
778. France 24 — *En cifras: 229 muertos y 17.800 millones de euros en pérdidas, a un año de la DANA* (October 28, 2025). https://www.france24.com/es/europa/20251028-en-cifras-229-muertos-y-17-800-millones-de-euros-en-p%C3%A9rdidas-a-un-a%C3%B1o-de-la-dana
779. Wikipedia — *Inundaciones de la DANA de 2024 en España* (237 dead in total: 229 in the province of Valencia, 7 in Castilla-La Mancha, 1 in Málaga; more than 2,600 injured; 11,242 homes; ≈141,000

vehicles; 800 km of roads). https://es.wikipedia.org/wiki/Inundaciones_de_la_DANA_de_2024_en_Espa%C3%B1a
780. Maldita.es — *Qué sabemos de los datos de la riada en la Rambla del Poyo durante la DANA* (November 7, 2024; the CHJ's emails of 13:42, 15:04, 16:13 and 18:43; the 17:00 and 18:00 flows; 2,282 m³/s at 18:55; the gauging station destroyed by the water toward 19:00, per the CHJ's November 4 statement). https://maldita.es/clima/20241107/datos-rambla-poyo-dana-valencia/
781. El Debate — *«A las 17.40 del día de la dana la CHJ registró que el caudal del barranco del Poyo era cinco veces superior al del Ebro»* (February 24, 2026). https://www.eldebate.com/espana/comunidad-valenciana/20260224/1740-dia-dana-chj-registro-caudal-barranco-poyo-cinco-veces-superior-ebro_388667.html
782. El Debate — *La Confederación del Júcar borra datos críticos del caudal del Poyo en plena crisis de la DANA* (November 5, 2024); esDiario — *Desaparecen y reaparecen los datos del barranco del Poyo de la web de la CHJ*. https://www.eldebate.com/espana/comunidad-valenciana/20241105/confederacion-jucar-borra-datos-criticos-caudal-poyo-plena-crisis-dana-valencia_242165.html · https://www.esdiario.com/comunidad-valenciana/241105/146265/desparecen-aparecen-datos-barranco-poyo-web-chj.html
783. Actualidad Valencia — *La CHJ omitió barrancos clave en su red de control y ocultó un medidor en Massanassa, según revela el SAIH*. https://actualidadvalencia.com/control-caudal-chj-barrancos-horteta-gallego-massanassa/
784. ORM — *La jueza de Catarroja llama a declarar como investigados en la causa por la DANA a Salomé Pradas y Emilio Argüeso* (investigation opened March 10, 2025; Investigating Court No. 3 of Catarroja, Judge Nuria Ruiz Tobarra). https://www.orm.es/informativos/la-jueza-de-catarroja-llama-a-declarar-como-investigados-en-la-causa-por-la-dana-a-salome-pradas-y-emilio-arg-eso/
785. El Independiente — *La causa de la DANA afronta su recta final tras el varapalo del TSJ a la jueza de Catarroja* (March 17, 2026; the High Court declines to investigate Mazón). https://www.elindependiente.com/espana/2026/03/17/causa-dana-afronta-recta-final-varapalo-tsj-jueza-catarroja/
786. The Objective — *La jueza de la dana une a la causa los vídeos del Cecopi y pide a À Punt la hora de cada uno* (October 2, 2025); El Español — *La jueza requiere a À Punt «todas las grabaciones de imagen y sonido» de la reunión del Cecopi* (the videos, from a production company hired by the emergencies secretariat; the originals demanded to fix each one's exact time). https://theobjective.com/espana/tribunales/2025-10-02/jueza-dana-videos-cecopi-punt-hora/ · https://www.elespanol.com/valencia/20250919/jueza-dana-requiere-punt-todas-grabaciones-imagen-sonido-reunion-cecopi-octubre/1003743931870_0.amp.html
787. Maldita.es — *«Nos están ocultando los muertos» del parking de Bonaire (Aldaia): un año después, cómo funcionó la maquinaria de la desinformación* (October 29, 2025; the false municipal report of "more than 220 corpses"; false documents with more than 800 victims; not one fatality in the garage). https://maldita.es/malditobulo/20251029/teoria-conspiracion-muertos-bonaire/
788. Newtral — *Un mes y más de 60 bulos desmentidos: la desinformación surgida a raíz de la dana* (November 29, 2024). https://www.newtral.es/desinformacion-dana/20241129/
789. COPE — *Quién está detrás de la creación de bulos de la DANA y qué herramientas digitales han usado* (generative AI as the main enabler); elDiario.es — *Los bulos y falsedades sobre la DANA de València* (interference with emergency work). https://www.cope.es/programas/la-linterna/audios/detras-creacion-bulos-dana-valencia-herramientas-digitales-han-usado-20241117_3046424.html · https://www.eldiario.es/comunitat-valenciana/bulos-falsedades-dana-valencia-cientos-cadaveres-bonaire-demolicion-pantanos_1_11791455.html
790. La Moncloa — *El Consorcio de Compensación de Seguros supera los 4.000 millones en pagos por indemnizaciones a afectados por la DANA* (November 26, 2025; 210,106 payouts). https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/economia-comercio-empresa/paginas/2025/261125-indemnizacion-dana.aspx
791. SegurosNews — *El Consorcio ya ha pagado 4.378 millones por la DANA de Valencia* (March 27, 2026; 212,237 claims paid; 251,249 files; 99.3% processed; total estimated cost ≈€4.8 billion). https://segurosnews.com/news/el-consorcio-ya-ha-pagado-4-378-millones-por-la-dana-de-valencia
792. Wikipedia — *Incendios forestales de España de 2025* (≈400,000 ha by August 19; Ourense ≈150,000 ha; Uña de Quintana ≈40,000 ha; A Rúa >38,000 ha; Molezuelas de la Carballeda, more than 700 evacuated; 500 additional UME soldiers on top of more than 1,400). https://es.wikipedia.org/wiki/Incendios_forestales_de_Espa%C3%B1a_de_2025

793. Vallfirest — *La oleada de incendios forestales de agosto de 2025* (a score of major fires in two weeks, more than 300,000 ha; the worst season in more than 30 years). https://www.vallfirest.com/noticias/oleada-incendios-forestales-2025
794. Lauegi Center of Aran (Conselh Generau d'Aran) — *Avalanche danger bulletin and ATES mapping* (daily bulletin; "the information does not replace training, experience and good judgment"). https://lauegi.conselharan.org/ates/?lang=en
795. ICGC — *Butlletí de Perill d'Allaus (BPA)* (the Catalan Pyrenees); AEMET — *Boletín de peligro de aludes* (the other ranges). https://www.icgc.cat/en/Thematic-areas/Risks-and-emergencies/Avalanches/Avalanche-Danger-Bulletin-BPA · https://www.aemet.es/en/eltiempo/prediccion/montana/boletin_peligro_aludes
796. Civil Protection (Spain) — *Public Warning System / ES-Alert* (Art. 110 of Directive (EU) 2018/1972, transposed by Law 11/2022 on Telecommunications; operational since June 2022 and available to the regional 112 centers since 2023; cell broadcast technology). https://www.proteccioncivil.es/coordinacion/redes/ran/public-warning-system · https://www.vademecumlegal.es/legislacion/articulo-110-codigo-europeo-comunicaciones-electronicas
797. esPúblico — *Responsabilidad patrimonial por daños por el estado de la calzada* (Art. 106.2 of the Constitution; Law 40/2015; the causal link; the Guardia Civil's report as the central evidence; the 6-month administrative silence). https://www.administracionpublica.com/responsabilidad-patrimonial-por-danos-por-el-estado-de-la-calzada/
798. t2p — product manual (SHA-512 fingerprint on the device or at the datum's source; RFC 3161/eIDAS timestamp; anchoring on the public blockchain; .asice container). See chapters 13 and 15.
799. Beth Simone Noveck — *Reboot: AI and the Race to Save Democracy* (Yale University Press, 2026), ch. 3 ("Ballots and Bots").
800. Munich Security Conference — *AI Elections Accord* (2024). https://securityconference.org/en/aielectionsaccord/
801. Utah State Legislature — *SB 131, Synthetic Media / disclosure of synthetic media* (2024). https://le.utah.gov/~2024/bills/sbillenr/SB0131.pdf
802. Microsoft / MSC — *Tech Accord to Combat Deceptive Use of AI in 2024 Elections* (16-02-2024). https://blogs.microsoft.com/on-the-issues/2024/02/16/ai-deepfakes-elections-munich-tech-accord/
803. Beth Simone Noveck — *Reboot* (2026), ch. 3 (Taiwan, Audrey Tang, Cofact).
804. The Washington Post — *AI mayoral candidate VIC, Cheyenne* (2024) https://www.washingtonpost.com/technology/2024/08/19/artificial-intelligence-mayor-cheyenne-vic/ ; Wikipedia — *Jason Palmer (politician)* https://en.wikipedia.org/wiki/Jason_Palmer_(politician ; TIME — *Albania's AI minister Diella* (2025) https://time.com/7324934/albania-ai-minister-diella/
805. Beth Simone Noveck — *Reboot* (2026), introduction and ch. 3 ("possibilist, not Pollyanna"; provenance "alone will not solve the problem").
806. Beth Simone Noveck — *Reboot* (2026), ch. 3 (the Korean media literacy and deepfakes trial).
807. IPI / VSquare / HKS Misinformation Review — *the Šimečka–Tódová audio deepfake in Slovakia* (28-09-2023; the 48-hour moratorium; the hole in Meta's audio policy; the debate over the real influence on the result). https://ipi.media/slovakia-deepfake-audio-of-dennik-n-journalist-offers-worrying-example-of-ai-abuse/ · https://vsquare.org/slovak-election-targeted-by-pro-kremlin-deepfake-hoax/ · https://misinforeview.hks.harvard.edu/article/beyond-the-deepfake-hype-ai-democracy-and-the-slovak-case/
808. Beth Simone Noveck — *Reboot: AI and the Race to Save Democracy* (Yale University Press, 2026), ch. 8; and *Post Office Horizon Inquiry* (Sir Wyn Williams, 2024). https://en.wikipedia.org/wiki/British_Post_Office_scandal
809. Beth Simone Noveck — *Reboot* (2026), ch. 8 ("Using procurement power").
810. Beth Simone Noveck — *Reboot* (2026), ch. 6 (Policy Synth, runtime guardrails); The GovLab. https://thegovlab.org/
811. Beth Simone Noveck — *Reboot* (2026), ch. 8 ("double opacity"; external audits).
812. Government of California — *Transparency in Frontier Artificial Intelligence Act (SB 53)* (29-09-2025). https://www.gov.ca.gov/2025/09/29/governor-newsom-signs-sb-53-advancing-californias-world-leading-artificial-intelligence-industry/
813. DeepSeek — *DeepSeek-R1 release* (20-01-2025). https://api-docs.deepseek.com/news/news250120

814. Beth Simone Noveck — *Reboot* (2026), ch. 8 (open source: the five elements; the record of queries and training runs).
815. Beth Simone Noveck — *Reboot* (2026), introduction and ch. 6 ("accountability in name only"; AI as an object).
816. Federal Senate of Brazil — *e-Cidadania portal / the Joca Law* (2024-2025) https://www12.senado.leg.br/ecidadania/ ; and *Reboot* (Noveck, 2026), ch. 7.
817. Audrey Tang and E. Glen Weyl — *Plurality: The Future of Collaborative Technology and Democracy* (2024); and *Reboot* (2026), ch. 7.
818. Beth Simone Noveck — *Reboot* (2026), ch. 3 (personal APIs).
819. Beth Simone Noveck — *Reboot* (2026), ch. 7 (scale, action, institutionalization; "participation without integration").
820. time2act — chapter 13 (t2p: the video-selfie signature and on-device proof of life); T2P-Manual (v1.0, 2026).
821. Beth Simone Noveck — *Reboot* (2026), ch. 3 (the Univ. of Madrid survey: 51% would replace politicians with AI).
822. NPR — *Pulitzer Board Recognizes Darnella Frazier For Recording George Floyd's Murder* (June 11, 2021; the special citation; the jury's verbatim wording). https://www.npr.org/2021/06/11/1005601724/darnella-frazier-teen-who-filmed-george-floyds-murder-wins-pulitzer-prize-citati
823. Wikipedia — *Darnella Frazier* (17 when she recorded; the video as evidence in the Derek Chauvin trial). https://en.wikipedia.org/wiki/Darnella_Frazier
824. elDiario.es — *Juicio a tres policías por acusar en falso a una manifestante de agresión: «Ella no está condenada porque había un vídeo»* (the "Rodea el Congreso" case; the initial report recorded only insults). https://www.eldiario.es/politica/juicio-tres-policias-acusar-falso-manifestante-si-no-llegamos-encontrar-video-hubieran-condenado_1_10922988.html
825. elDiario.es — *Absueltos los antidisturbios acusados de mentir en un juicio contra una manifestante del 'Rodea el Congreso'* (the acquittal; the court sees no crime despite the absence of evidence of the assault). https://www.eldiario.es/politica/absueltos-antidisturbios-acusados-mentir-juicio-manifestante-rodea-congreso_1_10979336.html
826. Togas.biz — *Si el atestado policial no ha sido ratificado en el plenario… las manifestaciones en él contenidas no tienen eficacia probatoria* (the doctrine on the evidentiary force of the police report). https://togas.biz/articulos/articulo-profesionales--si-el-atestado-policial-no-ha-sido-ratificado-en-el-acto-del-plenario-por-los-agentes-de-la-autoridad-que-lo-confeccionaron-someti-ndose-al-principio-de-publicidad-y-contradicci-n-las-manifestaciones-en-l-contenidas-no-tiene-la-eficacia-probatoria/
827. Infobae — *El artículo de la 'ley mordaza' que derogarán PSOE y Sumar fue declarado inconstitucional y los policías ya no lo aplican* (July 17, 2024; Art. 36.23; the Constitutional Court strikes down "unauthorized"; the agreed reform: a "demonstrable," reasoned risk). https://www.infobae.com/espana/2024/07/17/el-articulo-de-la-ley-mordaza-que-derogaran-psoe-y-sumar-fue-declarado-inconstitucional-y-los-policias-ya-no-lo-aplican/
828. Infobae — *Más de 400.000 multas y un tercer intento de reforma: la ley mordaza cumple 10 años* (March 31, 2025; fines under the law as a whole). https://www.infobae.com/espana/2025/03/31/mas-de-400000-multas-y-un-tercer-intento-de-reforma-la-ley-mordaza-cumple-10-anos-en-medio-de-las-criticas-sobre-su-impacto-en-derechos-y-libertades/
829. Amnesty International Spain — *Ley Mordaza: diez años de vulneraciones de derechos humanos y promesas incumplidas* (2025). https://www.es.amnesty.org/en-que-estamos/noticias/noticia/articulo/ley-mordaza-diez-anos-de-vulneraciones-de-derechos-humanos-y-promesas-incumplidas/
830. SME — Sindicat de Mossos d'Esquadra — *Consell de la Policia de 6 de febrer del 2025* (Axon Body 3 across 98% of the territory with 800 cameras; 5,000 officers trained; authorization for 7,000 and 1,500 more; the regulation being updated). https://sme-mossos.cat/consell-de-la-policia-de-6-de-febrer-del-2025/
831. USPAC — *Resum reunió subcomissió de materials, 27 de novembre de 2024* (300 more cameras planned; requests for generalization). https://uspac.cat/resums-reunions/resum-reunio-subcomissio-de-materials-27-novembre-2024/
832. Cámaras de Vigilancia IP — *Cámaras unipersonales de la Policía: funcionamiento y normativa* (the DGU, AB3 model; gradual extension; planned for the Guardia Civil; the aim: protecting the public and the officers). https://camarasdevigilanciaip.es/camaras-unipersonales-policia/

833. Telemundo — *Trump afirma falsamente que una foto de la multitud en un mitin de Harris en Detroit fue creada con inteligencia artificial* (August 2024; a photo by a staff member, unmanipulated; AP journalists on the tarmac; 15,000 attendees). https://www.telemundo.com/noticias/noticias-telemundo/elecciones-estados-unidos-2024/trump-afirma-falsamente-que-una-foto-de-la-multitud-en-un-mitin-de-har-rcna166423
834. Univision — *Esas fotografías de eventos de campaña de Kamala Harris no fueron creadas con IA, como afirman Trump y en redes* (the fact-check). https://www.univision.com/noticias/fotografias-eventos-campana-kamala-harris-no-ia-como-afirman-trump-y-redes
835. t2p — product manual (SHA-512 fingerprint on the device, by segments while recording; RFC 3161/eIDAS timestamp; anchoring on the public blockchain; .asice container; a continuous chain of fingerprints for body cameras; optional publication). See chapters 13 and 15.
836. Oxford Economics, *The global economic costs of climate inaction*. https://www.oxfordeconomics.com/resource/the-global-economic-costs-of-climate-inaction/
837. Climate Policy Initiative, *The Cost of Inaction*. https://www.climatepolicyinitiative.org/the-cost-of-inaction/
838. (a synthesis of climate-economic modeling) Earth.org. https://earth.org/have-economists-been-underestimating-the-cost-of-climate-change/
839. Orrick — *TAKE IT DOWN Act becomes law*; DEFIANCE Act (Senate, Jan-2026). https://www.orrick.com/en/Insights/2025/05/TAKE-IT-DOWN-Act-Becomes-Law
840. Beth Simone Noveck — *Reboot* (Yale UP, 2026), ch. 8; Post Office Horizon Inquiry (Sir Wyn Williams, 2024).
841. L2BEAT — *Decentralization status of Ethereum L2s (sequencers)* (2026). https://l2beat.com/scaling/summary — and MEV Watch, *blocks with OFAC censorship* (2026). https://www.mevwatch.info/
842. Jack Watling — *Statecraft: The New Rules of Power in a Divided World* (2023).
843. Dominic Frisby — *The Secret History of Gold* (2024).
844. Carissa Véliz — *Privacidad es poder* (Debate, 2021).
845. EDPB — *Opinion 08/2024 on "consent or pay"* (17-04-2024). https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-082024-valid-consent-context-consent-or_en
846. FTC — *Actions against data brokers (X-Mode/Outlogic) and Avast* (2024). https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-order-prohibits-data-broker-x-mode-social-outlogic-selling-sensitive-location-data — and Autoriteit Persoonsgegevens (NL), *€30.5 million fine on Clearview AI* (2024). https://www.edpb.europa.eu/news/national-news/2024/dutch-supervisory-authority-imposes-fine-clearview-because-illegal-data_en
847. Carissa Véliz — *The Ethics of Privacy and Surveillance* (Oxford University Press, 2024).
848. European Policy Centre — *Eurodac expansion (age 12→6; police access), 2024 Migration Pact*. https://www.epc.eu/publication/managing-migration-through-data-the-expansion-and-integration-of-eurodac/
849. Carissa Véliz (ed.) — *The Oxford Handbook of Digital Ethics* (Oxford University Press, 2022).
850. time2act — *Style guide (STYLE.md): "the promise and the peril"; "no techno-utopia without evidence"*.
851. CPJ (Committee to Protect Journalists) — *Israel-Gaza War* (263 journalists and media workers killed since 7-Oct-2023; 108 jailed; 174 wounded; "the deadliest and most deliberate effort to kill and silence journalists" the CPJ has ever documented) and *At least 5 Gazan journalists killed in Israeli strike on Nasser Hospital* (25-Aug-2025). https://cpj.org/issue/israel-gaza-war/ · https://cpj.org/2025/08/at-least-5-gazan-journalists-killed-in-israeli-strike-on-nasser-hospital/
852. Watson Institute (Brown University), Costs of War project — *News Graveyards: How Dangers to War Reporters Endanger the World* (Nick Turse, 1-Apr-2025: more journalists killed in Gaza than in the great wars of the twentieth century combined — a comparison the underlying database itself qualifies; ~13 Palestinian journalists killed per month); Foreign Policy — *War in Gaza the 'Worst Ever Conflict for Reporters'*. https://watson.brown.edu/costsofwar/papers/2025/Journalists · https://foreignpolicy.com/2025/04/01/gaza-war-worst-ever-for-reporters-costs-of-war-project/
853. Wikipedia (EN) — *2025 Nasser Hospital strikes* (25-Aug-2025; first strike ~10:00 against the position of the Reuters camera — which was broadcasting live — and second strike ~10 min later against rescuers and journalists; ≥22 dead, 5 journalists: Hussam al-Masri/Reuters, Mariam Dagga/AP, Mohammed Salama/Al Jazeera, Moaz Abu Taha, Ahmed Abu Aziz; "tragic mishap"); CNN — *New video reveals third strike* (27-Aug-2025).

https://en.wikipedia.org/wiki/2025_Nasser_Hospital_strikes · https://www.cnn.com/2025/08/27/middleeast/gaza-nasser-hospital-israel-attack-three-strikes-intl

854. CPJ — *Israel kills Al Jazeera journalists in targeted Gaza City airstrike* (Anas al-Sharif and four colleagues, 10-Aug-2025; the prior alert by the CPJ and UN rapporteur Irene Khan, Jul-2025); NPR — *Israeli airstrike kills Al Jazeera's Anas al-Sharif in Gaza*. https://cpj.org/2025/08/israel-kills-al-jazeera-journalists-in-targeted-gaza-city-airstrike/ · https://www.npr.org/2025/08/12/g-s1-82143/gaza-anas-al-sharif-al-jazeera-israel
855. Al Jazeera — *Family of Al Jazeera Gaza bureau head killed in Israeli air raid* (25-Oct-2023: Wael al-Dahdouh's wife, son Mahmoud, daughter Sham and grandson Adam) and *To kill a family: the loss of Wael Dahdouh's family*. https://www.aljazeera.com/news/2023/10/25/family-of-al-jazeera-gaza-bureau-head-killed-in-israeli-air-raid · https://www.aljazeera.com/news/2023/11/1/to-kill-a-family-the-loss-of-wael-dahdouhs-family-to-israeli-bombs
856. The Times of Israel / Al Jazeera / Courthouse News — the ban on international press access to Gaza: the FPA's petition rejected (Jan-2024, "no vested right"), new petition (Sep-2024), the government keeps the ban "for security reasons" despite the ceasefire (Jan-2026) and the Supreme Court postpones again (28-Jan-2026); more than two dozen news organizations demand access. https://www.timesofisrael.com/ban-on-media-access-to-gaza-should-stay-in-place-israel-tells-high-court/ · https://www.aljazeera.com/news/2026/1/7/media-body-condemns-israels-continued-ban-on-foreign-media-access-to-gaza · https://www.aljazeera.com/news/2026/1/28/israels-top-court-delays-gaza-press-access-ruling-amid-years-long-ban · https://www.courthousenews.com/over-two-dozen-global-news-organizations-call-on-israel-to-allow-gaza-access/
857. RSF (Reporters Without Borders) — *RSF files fifth complaint with ICC about Israeli war crimes against journalists in Gaza* (30-Sep-2025; five complaints since Oct-2023; ≥210 journalists killed, 56 with reasonable grounds to believe they were deliberately targeted; the "double bombing" pattern) and *"Gazawood": the deadly accusations against Gaza's journalists*. https://rsf.org/en/rsf-files-fifth-complaint-icc-about-israeli-war-crimes-against-journalists-gaza · https://rsf.org/en/gazawood-deadly-accusations-against-gaza-s-journalists
858. Al Jazeera / The National — *Palestinian journalist Bisan Owda and AJ+ win Emmy for Gaza war documentary* (25-Sep-2024, "It's Bisan From Gaza — and I'm Still Alive"; ~4 million followers; a failed discrediting campaign) and the photojournalist Motaz Azaiza. https://www.aljazeera.com/news/2024/9/26/palestinian-journalist-bisan-owda-and-aj-win-emmy-for-gaza-war-documentary · https://www.thenationalnews.com/arts-culture/film-tv/2024/09/26/bisan-owda-emmy-award-palestine/
859. Wikipedia (EN) / NPR — *Misinformation in the Gaza war* and *Pallywood* (a peak of 220,000 uses on X in Nov-2023 according to BBC Verify; the "doll" case published and retracted by the Jerusalem Post and amplified by Israel's official account on X; the false positive by the "AI or Not" detector on a real photo; Ron Dudai: documentation dismissed with a single word, "Fake"); NPR — *False 'crisis actor' claims spread about Israel-Hamas war*. https://en.wikipedia.org/wiki/Misinformation_in_the_Gaza_war · https://en.wikipedia.org/wiki/Pallywood · https://www.npr.org/2023/11/27/1214451419/civilian-deaths-are-being-dismissed-as-crisis-actors-in-gaza-and-israel
860. Human Rights Watch / CNN / Forensic Architecture — the explosion at al-Ahli Hospital (17-Oct-2023; 471 dead according to the Gaza Ministry of Health, 100-300 according to US intelligence) and the forensic divergence: HRW (pattern consistent with a propelled rocket; an independent investigation is needed), CNN and AP (a failed rocket from Gaza, most likely), NYT (cause uncertain), Forensic Architecture (projectile from the direction of Israel). https://www.hrw.org/news/2023/11/26/gaza-findings-october-17-al-ahli-hospital-explosion · https://www.cnn.com/2023/10/21/middleeast/cnn-investigates-forensic-analysis-gaza-hospital-blast · https://forensic-architecture.org/investigation/israeli-disinformation-al-ahli-hospital
861. OHCHR / CNN — the killing of Shireen Abu Akleh (11-May-2022): the UN's conclusion (Jun-2022: shots from Israeli forces, "seemingly well-aimed"), CNN's investigation (24-May-2022) and the Israeli army's admission (5-Sep-2022: "probably" its own fire; no charges). https://www.ohchr.org/en/press-briefing-notes/2022/06/killing-journalist-occupied-palestinian-territory · https://www.cnn.com/2022/05/24/middleeast/shireen-abu-akleh-jenin-killing-investigation-cmd-intl · https://www.cnn.com/2022/09/05/middleeast/idf-shireen-abu-akleh-investigation-intl/index.html
862. Forensic Architecture / Earshot — *The Killing of Hind Rajab* (29-Jan-2024: 355 impacts on the car; acoustic analysis of the call: shots from 13-23 m). https://forensic-

architecture.org/investigation/the-killing-of-hind-rajab · https://www.earshot.ngo/investigations/the-killing-of-layan-hamada-and-hind-rajab

863. ECCHR / Syria Justice & Accountability Centre / CBS — the Caesar files (~53,000 photos carried out of Syria in 2013; expert authentication of metadata and medical forensics) and the Koblenz verdict (13-Jan-2022: life imprisonment for Anwar Raslan for crimes against humanity). https://www.ecchr.eu/en/case/caesar-photos-document-systematic-torture/ · https://syriaaccountability.org/inside-the-raslan-trial-a-forensic-analysis-of-the-caesar-photos/ · https://www.cbsnews.com/news/syria-war-bashar-assad-police-officer-crimes-against-humanity-germany/
864. International Criminal Court — *ICC Pre-Trial Chamber I rejects the State of Israel's challenges to jurisdiction and issues warrants of arrest for Benjamin Netanyahu and Yoav Gallant* (21-Nov-2024; the war crime of starvation as a method of warfare; crimes against humanity; a warrant also for Mohammed Deif). https://www.icc-cpi.int/news/situation-state-palestine-icc-pre-trial-chamber-i-rejects-state-israels-challenges
865. International Court of Justice — *Application of the Genocide Convention (South Africa v. Israel)*, provisional measures of 26-Jan-2024, 28-Mar-2024 and 24-May-2024 (binding). https://www.icj-cij.org/case/192
866. Access Now (#KeepItOn) / EFF — *Israel must stop cutting Gaza off from the world* (19-Jun-2025: ≥23 communications blackouts since Oct-2023, from strikes on the fiber and lack of fuel; "civilian connectivity turned into a means and method of warfare") and *Connectivity is a Lifeline, Not a Luxury*. https://www.accessnow.org/press-release/keepiton-israel-must-stop-cutting-gaza-off/ · https://www.eff.org/deeplinks/2025/06/connectivity-lifeline-not-luxury-telecom-blackouts-gaza-threaten-lives-and-digital
867. NPR / GPS World / The Times of Israel — GPS jamming and spoofing in the region (2024: >50,000 flights affected; a ×6 increase; apps placing users in Beirut or Cairo; an Israeli air base identified as a source). https://www.npr.org/2024/04/22/1245847903/israel-gps-spoofing · https://www.gpsworld.com/israeli-air-base-identified-as-alleged-source-of-gps-disruptions-in-mideast/ · https://www.timesofisrael.com/gps-jamming-amid-wars-playing-havoc-with-airline-navigation-report/
868. Starling Lab / Content Authenticity Initiative — *78 Days* (Reuters + Starling, Nov-2020→Jan-2021: >200 photos of the presidential transition with cryptographic metadata anchored on Filecoin and Hedera; provenance verifiable by anyone). https://www.starlinglab.org/78days/ · https://contentauthenticity.org/blog/cai-spotlight-78-days-from-stanford-starling-lab-and-reuters
869. Guardian Project / WITNESS — *ProofMode* (since 2017: OpenPGP signature and fingerprint at capture, sensor snapshot, notarization — with OpenTimestamps among the supported systems; designed for citizen human rights documentation). https://guardianproject.info/apps/org.witness.proofmode/ · https://blog.witness.org/2017/04/proofmode-helping-prove-human-rights-abuses-world/
870. t2p — product manual (SHA-512 fingerprint computed on the device — only the fingerprint leaves the phone —; RFC 3161/eIDAS timestamp; anchoring on the public blockchain; forensic GPS with honest degradation; offline mode with deferred anchoring; commit-and-reveal; .asice package verifiable offline; certificate with QR). See chapters 13 and 15; commit-and-reveal, in chapters 45 and 46; the aid chain of custody, in chapter 50; selective disclosure and war-crimes evidence, in chapter 20; content credentials, in chapter 17.
871. Wikipedia (EN) — *Operation Metro Surge* (4-Dec-2025 → Feb-2026: "the largest immigration enforcement operation ever carried out" according to the DHS; thousands of ICE and CBP agents in Minneapolis–Saint Paul; ≥3,789 arrests, under 25% with criminal records; the general strike of 23-Jan-2026; Judge Schiltz: ≥96 court orders violated by ICE since 1-Jan; Judge Blackwell: the "overwhelming majority" of those arrested had legal presence); NPR — *Minnesota ICE crackdown will end* (12-Feb-2026: Tom Homan announces the end of the operation — which "has cost the lives of two American citizens" — without naming them); Sahan Journal — the deployment figures (more than 4,000 agents at the peak), revealed under court order. https://en.wikipedia.org/wiki/Operation_Metro_Surge · https://www.npr.org/2026/02/12/nx-s1-5712280/minnesota-ice-surge-ends · https://sahanjournal.com/immigration/hundreds-ice-agents-leaving-minnesota-operation-metro-surge/
872. Wikipedia (EN) — *Killing of Renée Good* (7-Jan-2026, Portland Avenue, Minneapolis: ICE agent Jonathan Ross fires three shots — windshield and driver's window — with the vehicle turning to move away, according to the New York Times reconstruction; in the car, her partner and a six-year-old child; the allegations by Noem and Trump; Ross's own recording, jerked upward at the critical

moment; the state investigation agency's access to the evidence, revoked the next day; Ross, three days off duty and reassigned out of state); The Intercept — *Video Shows ICE Agent's Fatal Shooting of Civilian in Minneapolis* (the close-range shot at the departing vehicle; a second angle shows Good yielding to the federal vehicles minutes earlier). https://en.wikipedia.org/wiki/Killing_of_Ren%C3%A9e_Good · https://theintercept.com/2026/01/07/video-ice-shooting-civilian-minneapolis/

873. Wikipedia (EN) — *Killing of Alex Pretti* (24-Jan-2026, Minneapolis: an ICU nurse at the veterans hospital, 37, gun permit and no record; the witnesses' videos, reviewed and verified by Reuters, BBC, Wall Street Journal, AP, CNN, New York Times and The Guardian: the phone in his hand, the pepper spray, the agents pinning him down, the gun pulled from his holster and the shots a second later; the DHS allegations — "maximum harm," "two magazines and no identification of any kind" — and Stephen Miller's; the Hennepin County medical examiner: homicide; the agents, identified by ProPublica); NPR — *Videos refute DHS account of fatal shooting* (25-Jan-2026: Noem — "approached with a gun" — against the videos: he never drew any weapon; he was filming with his phone); The Marshall Project — *How Many People Have Been Killed by ICE and Border Patrol?* (26-Jan-2026: no resuscitation by the agents; a pediatrician pushed away; in the Good case, three minutes without calling emergency services; the public rebuttal by Mayor Frey and Governor Walz after seeing the videos). https://en.wikipedia.org/wiki/Killing_of_Alex_Pretti · https://www.npr.org/2026/01/25/nx-s1-5687384/minneapolis-minnesota-shooting-video-dhs · https://www.themarshallproject.org/2026/01/26/ice-minneapolis-shooting-alex-pretti
874. ProPublica — *Minnesota's Fight to Hold Agents Who Shot Alex Pretti, Renee Good Accountable* (the lawsuit by Minnesota, Hennepin County and the state investigation agency against the Departments of Justice and Homeland Security, 24-Mar-2026: evidence withheld "to shield the agents involved" — names, weapons, reports, GPS data and witness statements denied despite formal demands; County Attorney Moriarty's seventeen criminal investigations; the supremacy-clause obstacle; the agents in the Pretti case, identified by ProPublica on 1-Feb); NPR — *Months after the ICE shootings in Minnesota, a federal probe remains elusive* (10-Apr-2026: Pretti's phone, still held; Good's car, sealed in an FBI warehouse unexamined; state investigators physically kept from the scene; the operation's only disciplinary measure: two agents on administrative leave for "false statements" in the case of the third person shot, Julio Sosa-Celis; "It looks like they are actively preventing the state from investigating" — Rachel Moran, law professor). https://www.propublica.org/article/minnesota-trump-ice-shooting-lawsuit-alex-pretti-renee-good · https://www.npr.org/2026/04/10/nx-s1-5775847/alex-pretti-renee-good-ice-shootings-federal-investigations
875. Wikipedia (EN) — *Shooting of Marimar Martinez* (Chicago, 4-Oct-2025: five shots by Border Patrol agent Charles Exum; the body camera shows the agents' vehicle turning into hers after "it's time to get aggressive"; the agent's messages — "fired 5 shots and she had 7 holes" —; the charges against Martinez and her co-defendant dismissed with prejudice on 20-Nov-2025) and *List of shootings by U.S. immigration agents in the second Trump administration* (the running count: 38 shootings and 9 dead between 20-Jan-2025 and 7-Jul-2026; the cases with charges against civilians dropped when video contradicted the agents' version); AP (via KSAT) — *Judge's blistering opinion details use of force in Chicago-area immigration crackdown* (Nov-2025: Judge Sara Ellis's 223-page opinion — the government's account, "simply not credible," contradicted by the body cameras she herself had ordered activated; Bovino's shifting versions about the stone and the gas; the use-of-force report drafted with ChatGPT). https://en.wikipedia.org/wiki/Shooting_of_Marimar_Martinez · https://en.wikipedia.org/wiki/List_of_shootings_by_U.S._immigration_agents_in_the_second_Trump_administrat · https://www.ksat.com/news/politics/2025/11/21/judges-blistering-opinion-details-use-of-force-in-chicago-area-immigration-crackdown/
876. KPRC 2 / Click2Houston — *ICE agent shoots, kills Lorenzo Salgado Araujo during attempted arrest in east Houston* (7-Jul-2026, Magnolia Park: ICE alleges he "rammed an official vehicle" and that the agent fired "in self-defense"; the agents wore no body cameras, according to the DHS; conflicting accounts of the preceding moments; the Harris County medical examiner, 9-Jul: homicide). https://www.click2houston.com/news/local/2026/07/07/heavy-law-enforcement-presence-reported-in-east-houston-neighborhood/
877. Leo XIV — *Encyclical letter "Magnifica Humanitas"* (Vatican, May 15-25, 2026; on the guardianship of the human person in the age of artificial intelligence). https://www.vatican.va/content/leo-xiv/es/encyclicals/documents/20260515-magnifica-humanitas.html

878. Leo XIII — *Rerum Novarum* (1891; the origin of the Church's social doctrine). https://www.vatican.va/content/leo-xiii/es/encyclicals/documents/hf_l-xiii_enc_15051891_rerum-novarum.html
879. Francis — *Laudato Si'* (2015; the common home and the technocratic paradigm). https://www.vatican.va/content/francesco/es/encyclicals/documents/papa-francesco_20150524_enciclica-laudato-si.html
880. Vatican News — *Leo XIV presents his first encyclical "Magnifica Humanitas"* (May 25, 2026). https://www.vaticannews.va/es/papa/news/2026-05/papa-leon-xiv-primera-enciclica-magnifica-humanitas.html
881. John Paul II — *Centesimus Annus* (1991; centenary of 'Rerum Novarum'). https://www.vatican.va/content/john-paul-ii/es/encyclicals/documents/hf_jp-ii_enc_01051991_centesimus-annus.html
882. Vatican News / The Washington Post — *Leo XIV presents "Magnifica Humanitas" in person, with Anthropic cofounder Chris Olah; "with our differences, we can listen to one another"* (25-05-2026). https://www.vaticannews.va/en/pope/news/2026-05/pope-leo-xiv-encyclical-magnifica-humanitas-ai.html
883. Anthropic — *Chris Olah's remarks on Pope Leo XIV's encyclical "Magnifica Humanitas"* (the Church as "informed critics"; "a long collaboration between those of us who are building this and those who can see what we, from inside, cannot"; dialogue with people "not motivated by profit") (2026). https://www.anthropic.com/news/chris-olah-pope-leo-encyclical
884. Fortune — *Who is Chris Olah? The atheist Anthropic cofounder the Pope chose to sit beside him at the Vatican* (raised evangelical in Toronto, an atheist since 15; the industry "can't govern itself") (03-06-2026). https://fortune.com/2026/06/03/who-chris-olah-anthropic-cofounder-atheist-pope-leo-vatican-ai/

List of figures

- Figure 0.1 — Timeline 2004 → 2026: the milestones in energy, climate, AI and blockchain that frame this book.
- Figure 0.2 — From "supply peak" to "peak demand": Roberts's projection (2004) versus the observed trajectory.
- Figure 1.1 — From centralized to distributed trust: the journey of a fact, from capture to a verification anyone can redo.
- Figure 2.1 — The world still runs on fossil fuels (~86.7% of primary supply in 2024), but renewables already lead the growth.
- Figure 3.1 — The price of Brent: from the 2022 shock (~$130) to stabilization near ~$67 in 2025.
- Figure 4.1 — The cost of solar electricity fell nearly 90% between 2010 and 2024: the fastest price transformation in the history of energy.
- Figure 4.2 — Electric vehicle sales pass 17 million in 2024 and already account for more than one in five new cars.
- Figure 5.1 — China dominates the manufacturing of the transition: from 60% of wind turbines to ~91% of rare-earth refining.
- Figure 6.1 — Global warming exceeded +1.5°C in a calendar year for the first time (2024).
- Figure 7.1 — Efficiency improvement (~1% in 2024) runs far below the pace needed to double it by 2030.
- Figure 8.1 — Despite the green rhetoric, 81% of the investment of the G20's state energy companies still goes to fossil fuels (2024).
- Figure 9.1 — Data center electricity consumption doubles by 2030 (~950 TWh), driven by AI.
- Figure 10.1 — Deepfakes multiply: from ~500,000 (2023) to 8 million (2025), on the way to 90% of online content in 2026.
- Figure 11.1 — AI orchestrates renewables, storage, demand response and new nuclear into a resilient grid.
- Figure 12.1 — Producer and consumer exchange energy and certificates directly; the blockchain provides trust without an intermediary.
- Figure 13.1 — From the hash to decentralized verification: no one has to trust anyone; only the mathematics.
- Figure 13.2 — C2PA proves "who and when" at origin; the blockchain anchor proves "this existed and has not been altered."
- Figure 13.3 — The t2p flow: capture, seal the fingerprint on the public blockchain, and verify it even without a connection or intermediaries.
- Figure 14.1 — Green hydrogen (€4-7/kg) is still far more expensive than gray or blue: the challenge is not technical, but one of cost and demand.
- Figure 15.1 — Detector accuracy decays in the real world as generators improve; a provenance anchor does not degrade with time.
- Figure 16.1 — In deepfake fraud, trust breaks at identity and instruction; a proof of existence anchored at the origin adds an independent verification point.
- Figure 17.1 — When an image is re-uploaded, the C2PA metadata can disappear; the hash anchored to the public chain still proves the original existed, unaltered.
- Figure 18.1 — If the image's hash is computed at the moment of acquisition and anchored with a timestamp, any later manipulation (an injected or erased tumor) is detected instantly.
- Figure 19.1 — The issuer signs and anchors the credential's hash; the holder presents it; any verifier checks its authenticity and integrity without intermediaries.
- Figure 20.1 — Capture hashed at the origin → metadata and geolocation → immutable anchoring → admissible evidence: the chain that turns digital testimony into proof.
- Figure 21.1 — The price of the European emission allowance (EUA) climbed to a peak of €105.73 in February 2023 and moved around €80 in mid-2026; in parallel, the CBAM entered its definitive phase on January 1, 2026, with a first certificate set at €75.36/t. Sources: Trading Economics; European Commission; S&P Global.
- Figure 22.1 — Across a credit's life — design and baseline, implementation in the field, MRV, issuance, sale and optional tokenization — fraud infiltrates above all through the inflated baseline and the manipulable field evidence. The proof of existence anchors the evidence links (photo, measurement, report) but does not fix a defective methodology. Sources: The Guardian; Verra; CarbonPlan; ICVCM.

- Figure 23.1 — From the measures catalog to certificate issuance: the measure is captured as evidence (before/after photo, meter, invoice), verified and converted into a CAE (1 CAE = 1 kWh). The proof of existence is inserted at the moment the evidence is captured — SHA-512, forensic GPS, RFC 3161 seal, liveness check and .asice container — exactly where French law already requires dated, geolocated photographs. Sources: BOE (RD 36/2023); Légifrance (Law 2025-594); Funcas.
- Figure 24.1 — Self-consumption with proof of existence (t2p: SHA-512 fingerprint on the device, RFC 3161/eIDAS, public blockchain, .asice): the inverter's and meter's own reading sealed in chained batches, the bill with its compensation cross-checked, the energy community's minutes (coefficients, joiners and leavers with a certain date) and the contract/PPA frozen on signing day. Context: 9.3 GW of self-consumption in Spain (UNEF) — fifth source by capacity; RD 244/2019 (simplified monthly compensation ≤100 kW; collective up to 5 km). Limits: one's own reading contrasts the fiscal meter, it does not replace it; calibration first; only fingerprints on the chain (the consumption curve is intimacy). Sources: UNEF; BOE (RD 244/2019).
- Figure 25.1 — A sensor reading (odometer, CAN bus, GPS, OCPP charging session) is digested into a SHA-512 fingerprint on the device, bound to forensic GPS and an RFC 3161/eIDAS timestamp with legal presumption, anchored to the public blockchain and packaged into an .asice container verifiable offline, ready to issue a CAE (sheets TRA010-TRA050), a carbon credit or a CSRD report. Comparison box: without a seal, up to 30-50% of cross-border odometers are tampered; with a sealed registry (the Car-Pass model), under 0.2%. Sources: European Parliament (EPRS); Car-Pass; MITECO.
- Figure 26.1 — The number of private cars a single shared vehicle removes from the street varies enormously by city and methodology: from an average of 3.3 in the European SHARE NOW study to the 7-11 of the car2go study in North America, with even higher scenarios in some cities. The publicity figure "one for ten" is commercial, not consensus. The lesson: never a single number; always a range with its source. Sources: UC Berkeley TSRC (Shaheen and Martin, 2016); Transportation Research Part A (2020).
- Figure 27.1 — The delay with proof of existence (sealed with t2p — chapter 13): the arrival with forensic GPS as the doors open (the CJEU's definition), the display board and flight status frozen, the trip file (ticket, boarding pass, care expenses) and the carrier's symmetry (anchored operating records; the extraordinary circumstance sealed as it happens). Context: Regulation (EC) 261/2004 (€250/400/600; Sturgeon ≥3 h); Renfe 1-7-2024: from 15/30 min to 60/90 (50%/100%). Limits: the judge decides what is extraordinary; document your own trip, never third parties. Sources: EUR-Lex; Renfe/Newtral.
- Figure 28.1 — Percentage of consumers willing to digitally alter the image or document of a claim, even against the insurer's rules: 36% on average, but up to 55% among Generation Z. The willingness falls with age. This is the face of the AI image fraud that 99% of insurers have already seen. Source: Verisk with the National Insurance Crime Bureau (2026).
- Figure 29.1 — At the point of origin (plot, harvest, factory) a georeferenced photo is captured and sealed with forensic GPS, an eIDAS timestamp and a SHA-512 fingerprint, and packaged into an .asice container verifiable offline. This proof-of-the-act-of-origin layer combines with serialization (the FMD's 2D code), the EUDR due diligence declaration or Aura's digital passport, through to final verification at customs or on the shelf. Sources: EUDR (EU 2023/1115); EMA; Aura Blockchain Consortium.
- Figure 30.1 — The consumer's information layer with proof of existence: the review with an anonymous attested purchase (the receipt sealed on the device; the review black market loses its raw material), the price series with a certain date (the lowest-price-in-30-days rule becomes verifiable) and the warranty clock (purchase, unboxing and notices sealed activate and attest the 3 legal years). Limits: the seal judges neither sincerity nor quality; purchase attestation must never identify the buyer. Sources: FTC (2024 rule); Amazon vs. Facebook groups; OCU (Black Friday); CEC/Consumo (3-year warranty).
- Figure 31.1 — Second-hand between individuals with proof of existence (sealed with t2p — chapter 13): the item filmed before shipping (overlaid date/time + forensic GPS), the sealing in a single take, the buyer's symmetrical opening and the dispute file (captured listing, chat, payment, tracking). Total symmetry: it disarms the fake package AND the "switcheroo." Context: >19 million Wallapop users; >€13 billion/year; second-hand scams as the bulk of the >412,000 computer frauds of 2024 (+60% complaints in 2025). Limits: proportionality by value; the hidden defect is for the expert; never film third parties. Sources: Guardia Civil/specialist press.
- Figure 32.1 — Renting things with proof of existence (sealed with t2p — chapter 13): the pickup walk-around filmed (overlaid date/time + forensic GPS of the parking lot), the return walk-around

— decisive in after-hours returns — the contract and the excess sealed, and the file against the late charge. Symmetry: the bilateral inspection protects the honest company from the customer who denies the dent; extendable to B2B machinery rental. Context: documented irregular practices (charges of €300-1,000 set unilaterally; inspections without the customer); the ECC already recommends photos — what was missing was the certain date. Limits: normal wear vs damage and prices are for the contract and the judge. Sources: Consumo/DSCA; CEC-España.

- Figure 33.1 — The customer before the company with proof of existence (sealed with t2p — chapter 13): the sealed cancellation (recording of one's own call, certified capture of the form), the contract and its versions (web capture with source code), the recorded verbal promise and the .asice file exportable to the consumer authority. Law 10/2025 grants the rights (3 minutes, tracking code, 15/5 days, simple cancellation); the certain date makes them enforceable — and vouches for the diligent company. Limits: one's own conversations, proportionality with the operator, the formal claim remains indispensable. Sources: BOE (Law 10/2025); FACUA (Vodafone/Movistar/Telefónica fines).
- Figure 34.1 — The collection clock with proof of existence (sealed with t2p — chapter 13): the delivery note photographed with forensic GPS at delivery, the invoice with a double certain date (issuance and dispatch), the automatic count of Law 3/2004 (30/60 days; late-payment interest 11.15% + €40) and the .asice file for the payment-order procedure. Context: VeriFactu (RD 1007/2023, postponed to 2027 by RD-Law 15/2025) already imposes invoicing records with chained hashes and QR — the state builds half a proof of existence; t2p adds the layer opposable between private parties. Limits: the seal neither collects invoices nor judges the work; VeriFactu protects the tax office, not the creditor. Sources: AEAT; BOE; PMcM 2025 (85% of large companies in breach; public administrations at 70 days).
- Figure 35.1 — The condition of the dwelling (rental) or a document (deed, contract) is captured and sealed on the device with SHA-512, forensic GPS, an eIDAS timestamp and proof of life, anchored to the public blockchain and packaged into an .asice container verifiable offline, ready for the registry, the notary, the court or the other party. The proof of existence supplies evidence with a certain date, not the legal title (which the registry creates). Context bar: ~70% of the world's population without registered title; ~1.1 billion people insecure about their property; Georgia, ~1.5 million titles anchored. Sources: World Bank; Prindex; Exonum/Bitfury.
- Figure 36.1 — The rental cycle with proof of existence: the accredited listing (title + seal, against the nonexistent-flat scam), the move-in inventory sealed bilaterally, the notices and repairs with a certain date during the contract, the chronology of possession (the tenant's contract and receipts; the owner's periodic condition checks; LO 1/2025 fast track) and the move-out inventory that settles the deposit without a war. Only the SHA-512 fingerprints are anchored to the public blockchain (RFC 3161/eIDAS, .asice); the images of the home never leave the devices. Sources: Guardia Civil (Op. ISLEDA); Idealista/Maldita; LO 1/2025.
- Figure 37.1 — The homeowners' association with proof of existence (sealed with t2p — chapter 13): the notice with a certain sending date, the minutes sealed the same day (the minute book as a chain of anchored versions), the traceable levy (quotes, agreement, invoices, certifications in one project) and reliable communication between residents and manager. Symmetry: it protects the challenging resident and the honest manager. Limits: meeting agreements to record, publication only by fingerprint (personal data), and no seal resolves the substance of disputes (Horizontal Property Act). Sources: Gijón conviction (Art. 253 CP); Horizontal Property Act.
- Figure 38.1 — The renovation with proof of existence (sealed with t2p — chapter 13): the quote and its versions sealed on the day of agreement, the initial condition photographed with forensic GPS (the evidence the works erase), the phases with milestones — paying against documented facts — and the acceptance signed with a video selfie by both parties. Symmetry: it protects the family from the advance scam and the honest professional from the client who will not pay. Context: Guardia Civil investigations (Vigo: ≥6 victims of up to €10,000; losses >€77,000 in one case). Limits: quality is for the experts; proportionality; never pay ahead of the facts. Sources: press/Guardia Civil; consumer law.
- Figure 39.1 — The two reference projections for the tokenized RWA market in 2030 differ about eightfold: McKinsey's prudent scenario (~$2 trillion, range 1-4) and the optimistic BCG/ADDX one (~$16.1 trillion); the real starting point was about $24 billion in mid-2025. The distance depends, at bottom, on how much "garbage" is admitted: on the gap between the token and the proof of its backing — the oracle problem — which the proof of existence helps to close. Sources: McKinsey (2024); BCG/ADDX (2022); CoinDesk/RWA.xyz (2025).

- Figure 40.1 — At the moment of creation (t0), the work is sealed with a SHA-512 fingerprint on the device, an anchor to the public blockchain, an eIDAS timestamp, GPS and proof of life: a certain date of what, when, where and who. On publication (t1), the C2PA credentials can be erased on upload to the networks; when a model ingests it or someone copies it (t2), the .asice container still proves the file existed intact at t0, before t2. The proof of existence proves precedence, not originality. Sources: C2PA; U.S. Copyright Office; the NYT, Getty and Andersen cases.
- Figure 41.1 — The inventor seals the disclosure document with a SHA-512 fingerprint, an anchor to the public blockchain and an eIDAS timestamp (certain date, integrity and authorship), at a marginal cost against the $15,000-25,000 of a patent. If they also publish it (Technical Disclosure Commons, IP.com, the web), it becomes prior art indexed by Espacenet or Google Patents; when a rival's application arrives, the examiner finds the dated disclosure and rejects it for lack of novelty. The seal proves date and content, not patentability — which the examiner decides. Sources: USPTO; WIPO; Technical Disclosure Commons.
- Figure 42.1 — The trade secret with proof of existence (t2p: SHA-512 fingerprint on the device — the content never leaves — RFC 3161/eIDAS, public blockchain, .asice): the sealed deposit (possession + perimeter + the reasonable measure of Law 1/2019's third requirement), the know-how's version-by-version chronology, the employee departure record and the NDA with a fingerprint perimeter and video-selfie signature. Symmetry: the innocent defendant's chain of dates proves independent development. Mirror of chapter 41: the patent protects by publishing; the secret, by concealing — which is why only the proof that does not reveal can defend it. Sources: the Waymo/Uber/Levandowski case; Law 1/2019.
- Figure 43.1 — Every step of public money — the bid or tender document, the recorded award and the delivered work (photo with GPS and date) — is digested into a SHA-512 fingerprint, anchored to the public blockchain and given an RFC 3161/eIDAS timestamp, with proof of life at the work site; everything is packaged into an .asice container that a citizen or an auditor can verify independently and offline. Context: world public procurement ≈$13 trillion a year; between 10% and 30% of public works lost to corruption and mismanagement. The proof of existence demonstrates integrity, moment, place and person — not the honesty of the decision or the quality of the work. Sources: World Bank; CoST; OECD; eIDAS Art. 41.
- Figure 44.1 — The citizen-administration relationship with proof of existence (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, public blockchain, .asice): the filing with sealed content (what and when), the clock of administrative silence with dated deadlines (Law 39/2015: 3 months; certificate in 15 days), the notification with a neutral trail (the 10 days auditable by both parties) and auditable appointment availability (against the appointment black market). Symmetry: the same trail vouches for the administration's diligence. Limits: it creates no appointments and resolves no files; no central citizen profile. Sources: Law 39/2015; Ombudsman (immigration appointments).
- Figure 45.1 — The draw no one can load, in two halves: the prior commitment (participant list and allocation rule sealed with t2p — SHA-512 fingerprint, RFC 3161/eIDAS, public blockchain — and published by fingerprint BEFORE the draw) and verifiable public randomness (the drand beacon of the League of Entropy: unbiased, unpredictable, verifiable values generated by threshold cryptography among independent operators). Result: an allocation recomputable by anyone — result = rule(list, randomness). Limits: a clean draw does not cure scarcity; the public ritual and the notary are reinforced, not replaced. Sources: examination and subsidized-housing draws; drand/League of Entropy.
- Figure 46.1 — The test sealed beforehand (t2p: SHA-512 fingerprint — only the fingerprint is anchored, the content never leaves custody — RFC 3161/eIDAS, public blockchain, .asice): the exam certified before the sitting with logged accesses, the marking key committed before the exam (annulments transparent and dated), the candidate's answer sheet dated on the way out, and the verifiable challenge file. Symmetry: the clean tribunal can prove it touched nothing — the Osakidetza case (19 exams with signs of leakage, resignations, eight years and shelving for lack of evidence) resolved in days, in either direction. Pattern: chapter 45's prior commitment, applied to the content. Sources: Redacción Médica; Deia; elDiario.es.
- Figure 47.1 — A person without papers generates a proof-of-life gesture ("it's me, I'm alive"); on the device its SHA-512 fingerprint is computed with forensic GPS and an eIDAS timestamp ("here, now"), anchored to the public blockchain, and can accompany a W3C verifiable credential; the holder keeps the proof on their device and shares it with selective disclosure, in an .asice container verifiable offline. Scale of the challenge: ~800 M without identity, ~4.4 M (up to 10 M) stateless, 150 million unregistered children. Sources: World Bank ID4D; UNHCR; UNICEF; W3C.

- Figure 48.1 — At the station, the tally sheet is photographed and sealed with a SHA-512 fingerprint, an eIDAS timestamp, forensic GPS and the official's proof of life, in an .asice container anchored to the public blockchain; if the figure is altered in transit to the counting center, the fingerprint does not match and exposes it. What the proof of existence does NOT do: it is not internet voting, it is not blockchain voting, and it does not prove the count is true — that is audited on the paper (RLA). Sources: Supreme Court of Kenya (2017); Rivest et al. (2021); Philip Stark.
- Figure 49.1 — At the point of delivery, a delivery receipt (anonymous beneficiary, amount, place, time) is sealed with a SHA-512 fingerprint, an eIDAS timestamp and forensic GPS, in an .asice container verifiable offline; the on-device proof of life confirms a living person without a central scanner; and Janus-style deduplication reveals only one bit ("already registered" or "new"), creating no biometric database. Real impact: in Ukraine ~$35 million of duplicated aid was avoided. Limit: it does not prove what happens to the cash once collected (Yemen, Congo). Sources: WFP; EPFL/ICRC (Janus); Human Rights Watch.
- Figure 50.1 — The chain of custody of relief with proof of existence (t2p: SHA-512 fingerprint on the device, RFC 3161/eIDAS seal, public blockchain, .asice): the manifest sealed at origin (scanner of up to 50 pages), the customs and warehouse certificates with forensic GPS (genuine fix ~50 m / ≤30 s; a simulated location blocks certification), the convoy with captures resumed when coverage returns, and the receipt signed with a video selfie and liveness challenges — exportable as an .asice package verifiable offline by the donor's auditor. Limits: the seal does not attest to the quality of the contents, does not lift blockades and never files the victims. Sources: IFRC (self-disclosed Ebola fraud); UN/WFP (Somalia 2010); OFAC (CLAP network); 2026 Venezuela earthquake.
- Figure 51.1 — Three anchored chronologies checked against each other (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, public blockchain, .asice): the sealed trail of the one who flees (survives the confiscation of the phone; CJEU C-136/24 P), the observers' documentation (testimonies with a certain date; Black Book of Pushbacks) and the authorities' anchored records (logbooks and positions; the Pylos case). The golden rule: anchored evidence adds credibility; its absence never subtracts. Limits: no registry of people, physical safety first, and the seal decides neither asylum nor migration policy. Sources: HRW/Amnesty (Pylos); RSA (indictment of 17 coast guard members); BVMN; CJEU C-136/24 P.
- Figure 52.1 — Three inputs converge on an incontestable will: the document, whose SHA-512 fingerprint is anchored to the public blockchain (integrity); the proof of life on video, which binds identity, presence and volition (against undue influence and forgery); and the eIDAS timestamp with GPS, which fixes which version is definitive (against "a later will exists"). The result is an .asice verifiable offline. Limit: it does not replace witnesses, notary or probate, and capacity is decided by the judge or the doctor. Sources: Cerulli Associates; e-Estonia/Guardtime; the Voysey case.
- Figure 53.1 — Co-parenting logistics with proof of existence (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, public blockchain, .asice), always bilateral and agreed: communication with a certain date, the handover attested from both devices (never cameras or photographed minors), expenses with anchored receipts and the versioned parenting plan. The compass: the best interests of the child; the record is a shield, not a weapon. Limits: it does not decide custody, does not measure parenting, and in situations of violence the priority is the specialized channels. Sources: INE (joint custody 49.7%, 2024); OurFamilyWizard (courts of 50 states); the parenting plan.
- Figure 54.1 — The exporter issues a bill of lading that is sealed with a SHA-512 fingerprint, an RFC 3161/eIDAS timestamp and GPS, anchored to the public blockchain (.asice). Bank A finances it and marks it as pledged. When a fraudster presents the "same" document to bank B, the fingerprint matches and shows it as already pledged: the second financing is blocked. Limit: the seal does not prove the merchandise physically exists in the warehouse (Qingdao's "ghost metal"). Context: $2.5 trillion trade finance gap; real electronic bills of lading, only 5.7%. Sources: Asian Development Bank; ICC; the Qingdao case.
- Figure 55.1 — In the vulnerable traditional model, the bank statement passes through the client (Wirecard) before reaching the auditor: that is the point where the forgery slips in. With t2p, the bank confirmation, the general ledger or the invoice is sealed at the source — SHA-512 fingerprint, eIDAS timestamp, GPS, anchoring to the public blockchain — and the auditor verifies an unalterable chain of custody (.asice). Limit: the seal does not make true a datum fabricated at the source (Wirecard, Luckin, Enron). Context: Wirecard, a nonexistent €1.9 billion; financial statement fraud, a median of $766,000 per case. Sources: Financial Times; PCAOB; ACFE.
- Figure 56.1 — Of every dollar an advertiser puts into programmatic buying, only 36 cents reach the consumer: 29 are lost in transaction costs and 35 in "media productivity loss" (junk websites and

low-quality inventory). Digital ad fraud cost $84 billion in 2023 — 22% of online spend — and bots are already half of internet traffic. Sources: ANA (2023); Juniper Research; Imperva.

- Figure 57.1 — System events (including Event ID 1102, "the log was cleared") are forwarded in real time off the attacked system; the log batches are summarized in a Merkle root that is sealed with a SHA-512 fingerprint and an RFC 3161/eIDAS timestamp and anchored to the public blockchain, chaining each root to the previous one (append-only). Any later alteration breaks the fingerprint and exposes the attacker. Limit: the anchor does not prevent the breach; it only protects what was recorded and forwarded in time. Context: average dwell time of 14 days in 2025; average cost of a breach, $4.44 million. Sources: MITRE ATT&CK; Mandiant M-Trends; NIST SP 800-86.
- Figure 58.1 — The breach seen from the citizen's side, with proof of existence (sealed with t2p — chapter 13): the notification captured (or its absence, dated), the certificate of presence in the leak (web capture with source code), the diary of the harm (phishing, calls, attempts — each piece sealed in the moment) and the .asice claim. Foundation: CJEU C-300/21 — non-material harm compensable with no threshold, but with infringement + harm + causation PROVEN. Symmetry: the diligent company proves its notification. Limits: valuation belongs to the judge; correlation ≠ causation; the diary stays on the device. Sources: CJEU; GDPR Art. 82.
- Figure 59.1 — The documented outage with proof of existence (sealed with t2p — chapter 13): the independent probe with its chain of sealed logs (the customer's clock), the provider's status page frozen with source code, the business impact documented (screen recording, metrics by fingerprint) and the .asice claim — for the SLA credits, the insurer and the DORA/NIS2 reports. Symmetry: the provider that anchors its availability makes it its commercial argument. Limits: the probe attests to when it was not responding, not the cause; methodology first; no seal replaces redundancy. Sources: CrowdStrike 19-7-2024 (8.5 million Windows machines; Delta ~$550 million).
- Figure 60.1 — On a legal dig (left), the piece is photographed in situ with GPS coordinates and an eIDAS timestamp before being moved: the provenance is documented from second zero and anchored to the public blockchain (.asice). In plunder (right), the paper "provenance" begins years after the find and is easy to fabricate — like the Chola Shiva sold to Australia for $5.6 million with false papers. Limit: the seal proves integrity, date, place and person, but does NOT authenticate the object (a sealed photo of a forgery is still a forgery). Context: illicit antiquities trade of $225 million to $6 billion/year. Sources: INTERPOL; Manhattan district attorney's office; Art Loss Register.
- Figure 61.1 — Four fronts of one defense (SHA-512 fingerprint, RFC 3161/eIDAS timestamp, anchoring to the public blockchain, .asice): the anchored manifests of archives and newspaper collections (silent edits detectable), survivors' testimonies sealed in life (temporal priority against fabrication), exhumation files with an anchored chain of custody (Law 14/2022, State DNA Bank) and the pre-AI horizon of photographic collections (the verifiable canon against synthetic history). Limits: the seal does not do source criticism or freeze interpretation; plurality of archives, never a single "archive of truth." Sources: UNESCO/WJC; Maldita.es; Internet Archive (2024 breach); Law 14/2022.
- Figure 62.1 — The whistleblower computes the document's SHA-512 fingerprint on their device and anchors it to the public blockchain with an eIDAS timestamp: WHAT and WHEN are proven, not WHO. The layer that would betray them — identity, GPS, proof of life, metadata, IP — is deliberately omitted: what protects is what is not sealed. With a zero-knowledge proof, the journalist or the court verifies that the document is authentic and predates a certain date, without knowing the source. Limit: the chain does not block printer tracking dots (the Winner case), stylometry or network metadata. Sources: ICIJ (Panama Papers); EFF; Freedom of the Press Foundation.
- Figure 63.1 — Across the life of an experiment, fraud can enter at the design (outcome switching: COMPare found that only 76.3% of primary outcomes were reported properly), at the sample (fabrication at origin, like Stapel), at the capture of the raw file or at the editing of the figure (band duplication: Bik found it in 3.8% of 20,621 papers). The seal at acquisition — SHA-512 fingerprint + eIDAS timestamp — covers the preregistration of the protocol and the capture of the raw datum, and makes any later edit detectable; but it does not prove that a datum fabricated at origin is true. Sources: COMPare (Oxford); Bik et al. (mBio 2016); ASM/ImageTwin.
- Figure 64.1 — Each attestation cycle starts from a fresh challenge of the day (nonce); the pensioner performs the liveness gesture on the device; the SHA-512 fingerprint of {liveness + challenge + GPS + time} is computed locally, anchored to the public blockchain with an eIDAS seal and closed in an .asice verifiable offline. The unrepeatable challenge prevents reusing old recordings: the dead cannot answer today's challenge. The face never leaves the device; only a fingerprint travels on the chain. The cost of not knowing who is still alive: Japan, 234,354 unconfirmed centenarians; U.S., $327

million to the deceased; South Africa, 140 million rand to 75,000 dead. Sources: France 24; SSA OIG; SASSA/Moneyweb.

- Figure 65.1 — Proof of life in the deepfake era: against virtual kidnapping with a cloned voice and AI-fabricated "proof of life" (FBI/IC3 alert, 2025), the three-piece protocol — unpredictable challenge, freshness of the moment and the t2p seal (SHA-512 fingerprint on the device, RFC 3161/eIDAS, public blockchain, .asice). Applications: the family protocol, the negotiation with dating, the voluntary last point on the map and the verifiable return (CNDES: 132 active alerts in 2025). Limits: the seal attests to freshness, not freedom; voluntary or nothing; never a substitute for police and negotiators. Sources: CNN (DeStefano case); FBI/IC3; Ministry of the Interior/CNDES.
- Figure 66.1 — The life of an anti-doping sample passes through five nodes: collection and the athlete's signature, sealing of bottles A and B, transport, reception at the laboratory and analysis. In Sochi, the fraud entered at the laboratory node (the "mouse hole" through which the secret service swapped urine at night). The t2p seal — forensic GPS, eIDAS timestamp and the kit number's fingerprint anchored to the public blockchain — makes that swap leave a detectable temporal and geographic discontinuity. The forensic evidence that exposed the fraud despite the "tamper-proof" bottles: 44 caps with tampering marks (12 from medalists), 8 samples with impossible salt, 2 with male DNA, ≥643 "disappeared" positives. Limit: the seal does not prove what was in the bottle before it was sealed. Sources: McLaren Report; CNN; Chemistry World.
- Figure 67.1 — Streaming music distributes a fixed-sum pool: in 2025, more than $22 billion (IFPI). Each fake stream invents no new money; it steals its share from artists with real audiences. The industry estimate puts the leak from fraud at some $2 billion a year — at least 10% of activity (Beatdapp). The Smith case: 661,440 bot streams a day, at half a cent, for seven years. With proof of life, listens without a human do not enter the distribution and the pool returns to real artists. Sources: IFPI Global Music Report 2026; Beatdapp; the New York prosecutors' indictment.
- Figure 68.1 — The network layer (STIR/SHAKEN) authenticates the originating number — "from what number is the call coming?" — but only 42.5% of calls arrive signed (April 2025), because any non-IP segment breaks the signature. The endpoint layer (t2p: proof of life + SHA-512 + eIDAS seal + GPS) authenticates the human origin and the moment of the act — "who, what and when, really?". They are complementary. Context: 52.8 billion robocalls in the U.S. in 2024; $48.8 million stolen via SIM swap (FBI, 2023). Sources: TransNexus; FCC; FBI IC3.
- Figure 69.1 — The "authorized" payment fraud with proof of existence (sealed with t2p — chapter 13): the deception documented (own call recorded, texts and chats captured), the session captured (the screen with the warning — or without it), the dated claim (deadlines) and the diligent bank's symmetry (warnings and frictions anchored in batches; the receiving account's onboarding verified). Context: United Kingdom, the world's first mandatory reimbursement (7-10-2024; £85,000; 50/50 sender-receiver; 5/35 days); PSD2: in unauthorized transactions, the burden of proof is the bank's. Limits: gross negligence is for the judge to weigh; prevention first. Sources: PSR; PSD2/RDL 19/2018.
- Figure 70.1 — Every custody event of a defense-grade component — marking by the original manufacturer, reception at the distributor, "origin" declared by the broker, AS6171 test report, contractor acceptance, installation — can be sealed with a SHA-512 fingerprint, an eIDAS timestamp, forensic GPS and proof of life, anchored to the public blockchain. The blind spot (in red): the physical substitution of the part *after* testing, which the seal does not detect without an anchor to the silicon (DNA, physically unclonable fingerprint). Context: ~80% of fake chips are recycled or remarked; the GAO bought 16 military parts online and all 16 were fake; a single broker introduced 84,000. Sources: U.S. Senate (2012); GAO; IHS/ERAI.
- Figure 71.1 — The software supply chain with proof of existence: every published version is anchored (SHA-512 fingerprint of the tarball, RFC 3161/eIDAS timestamp, public blockchain, .asice) and checked against the source code (the xz-utils discrepancy, detectable); the publishing chain watched in the transparency logs (Sigstore/Rekor: >101 million entries — proof of existence in production) triggers the alarm on releases with stolen keys (the Shai-Hulud worm); the inventory (SBOM) sealed per version makes incident response a query (Regulation (EU) 2024/2847). Limits: the seal attests to origin and integrity, not the code's goodness; maintainers need free tools and funding. Sources: CVE-2024-3094; CISA (Shai-Hulud); Sigstore/Rekor; CRA.
- Figure 72.1 — Fraud and opacity by vertical: three out of four industrial fishing vessels emit no public position signal (75%, Global Fishing Watch); a third of U.S. seafood is mislabeled — up to 87% of "red snapper" (Oceana); a global review across more than thirty countries finds 36% mislabeling; and up to 90% of "labeled clean" mineral in two zones of South Kivu came from illicit sources (Global Witness). t2p seals the origin at the point of extraction, before mixing or relabeling;

what the seal does NOT validate is the species or the legality of the zone. Sources: Global Fishing Watch; Oceana; Global Witness.

- Figure 73.1 — The journey of a donated dollar, with its real leaks: the professional fundraiser keeps 25 cents on average (New York attorney general, 2024), and in half the campaigns the charity receives less than half; in the Reynolds fraud, less than 3 cents reached the patient; in the Haiti project, 40% was lost to overhead. As a counterpoint, an effective NGO averts a death for about $3,000-5,500 (GiveWell). t2p seals three nodes — the restricted promise, every disbursement and the evidence of impact (GPS + time + proof of life) — turning "trust the report" into auditable evidence. Limit: it does not judge whether the declared metric is honest or the program's effectiveness. Sources: New York attorney general; FTC; ProPublica; GiveWell.
- Figure 74.1 — The recorders (the aircraft black box with more than 88 parameters and now up to 25 h of voice; the car's recorder with the 20 s before the crash; a drone's flight log) capture the data. t2p does not improve the capture: at the instant the recorder closes an event, it computes its SHA-512 fingerprint on the device and anchors it to the public blockchain with an eIDAS seal, so that denying the data (Tesla), trimming it (Cruise) or altering it leaves a fingerprint that does not add up. Limit: the seal does not correct an erroneous datum, does not settle the expert dispute over the cause (EgyptAir 990), and does not fill in what the sensor did not record (Jeju Air, where the recording cut out 4 minutes early). Sources: NTSB; California DMV; Benavides v. Tesla.
- Figure 75.1 — The personal black box: the audio of one's own conversation (STC 114/1984), the dashcam clip (AEPD limits) or the dated photo are encrypted and NEVER leave the device; only the SHA-512 fingerprint with the RFC 3161/eIDAS timestamp is anchored to the public blockchain (.asice). Disclosure is the holder's decision, before the judge. Two symmetrical outcomes: the victim who proves the harm and the innocent who proves the alibi. Limits: it does not attest the truth of the content, does not legalize recording a third party (Art. 197 CP), requires revocability and no central registry. Sources: STC 114/1984; State Prosecutor General's Office (Cometa; false reports 0.0082%); AEPD; UNESPA.
- Figure 76.1 — Two truths of genetic data. On the left, the family effect: it is enough for 2% of a population to be in a database to reidentify almost anyone through a distant relative, and close to 60% of those of European descent are already identifiable even if they never took a test (Erlich et al., 2018). On the right, the breach's multiplier effect: the attack on 14,000 accounts (0.1%) exposed 6.9 million users through the relative-matching feature. t2p can seal the act of consent — what, when and where — but it cannot restore privacy to an irrevocable datum that also belongs to your family. Sources: Science (Erlich et al.); TechCrunch; New York attorney general.
- Figure 77.1 — Digital evidence suffers two opposite attacks. On the left, the fake that deceives: an AI-fabricated audio or video presented as authentic (like the edited audio in a British family case). On the right, the liar's dividend: an authentic recording denied as a "deepfake" to dodge accountability (as in Huang v. Tesla, where it was argued that the Musk videos "could be deepfakes"). The same defense works for both: sealing the file in the second of recording with a SHA-512 fingerprint, an eIDAS seal and GPS (compatible with rule 902(14): the hash self-authenticates copies). Limit: the seal proves integrity and moment, not the truth of the scene — a sealed camera that records a setup certifies the setup. Sources: Chesney and Citron (2019); ProPublica; the Tesla and Baltimore cases.
- Figure 78.1 — A clinical note seen in two lanes. Above, the printed version: a single point, "02:03 — medication note, Dr. X." Below, the real audit trail with the t2p seal: the creation event declared at 02:03 but with the fingerprint anchored to the public blockchain at 08:11 — a six-hour gap — made from a non-clinical workstation (forensic GPS outside the care area), with a later backdated modification that cannot obscure the previous information (21 CFR 11.10(e)). The paper shows a point; the seal shows the story. Limit: the seal proves who and when, not the clinical truth of what was written. Sources: HIPAA and FDA (audit controls); Prieto v. Rush; forensic analysis of records.
- Figure 79.1 — A model's life cycle in four links: training data → training → weights + AI-BOM (model card, evaluation report) → deployment. t2p anchors the SHA-512 fingerprint of the weights and the manifest to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). Two risk points: illegal or opaque data at the origin (LAION-5B) and weights manipulated between the laboratory and the user (*nullifAI* backdoor via *pickle*), which the fingerprint check detects. Limit: the seal proves what was declared and when, and that the weights are the evaluated ones; it does not prove the declaration is true or the model safe. Sources: Stanford Internet Observatory (LAION-5B); JFrog (malicious models); European AI Regulation, Arts. 53 and 50; CycloneDX ML-BOM.
- Figure 80.1 — The verifiable mandate for AI agents with proof of existence (sealed with t2p — chapter 13): the mandate sealed before activation (budget, scopes, vetoes), the action log anchored in chained batches, the certified capture of the offer the agent saw (image + source code) and the

strong authorization with video selfie above the threshold. Context: the industry is already building signed mandates (Visa Intelligent Commerce, Mastercard Agent Pay, Google AP2 — Intent/Cart Mandate as verifiable credentials); the missing layer is the neutral one between user, platform and merchant. Precedent: Moffatt v. Air Canada (2024) — whoever deploys the agent answers. Sources: British Columbia tribunal; Visa/Mastercard/Google AP2.

- Figure 81.1 — Two paths. Above: the person creates, with proof of life, a bounded consent grant (which uses, for whom, until when) bound to a biometric reference (SHA-512 fingerprint; the face does not leave the device), anchored to the public blockchain, with a complementary eIDAS seal (.asice); a legitimate replica carries a verifiable pointer, and the revocation is dated. Below: a deepfake ad with no sealed grant — the absence itself proves the lack of consent. Limit: the seal proves whether there was consent, from whom and when; it does not detect the deepfake or prevent its creation, and it does not judge whether the consent was free. Sources: the Tom Hanks case; the Johansson–OpenAI case; the SAG-AFTRA agreement; the ELVIS Act; the Danish reform.
- Figure 82.1 — The minor's image with proof of existence (sealed with t2p — chapter 13) — the decisions of the ADULTS are documented, never the minor: the precise, sealed school authorization (uses, channels, duration), the dated disagreement between parents, the takedown request with dated response or silence, and the fingerprint of the post for the future (the grown child will be able to attest what was published about them, without any content traveling). Context: a child appears in ~1,300 photos online before age 13; French Law 2024-120 (the image right within parental authority, exercised jointly; judicial delegation in grave cases). Limits: the child's best interests are set by the judge; proportionality; prudence first. Sources: Légifrance; DAJ.
- Figure 83.1 — The victim or the support organization computes on the device the SHA-512 fingerprint of the content (the image does not leave the phone) and adds where and when it was found; the whole's fingerprint is anchored to the public blockchain, with a complementary eIDAS seal (.asice). Two outcomes: on the left, the provable takedown across all platforms (the StopNCII model, with a sealed 48-hour receipt); on the right, dated, admissible evidence for the court (as in Almendralejo). Limit: the seal does not remove the image, does not prevent its creation and does not detect the deepfake; never a registry of images, only fingerprints. Sources: the Almendralejo case; the Taylor Swift case; the TAKE IT DOWN Act; Directive (EU) 2024/1385; StopNCII.org.
- Figure 84.1 — The victim, the family or the school computes on the device the SHA-512 fingerprint of the hateful content (the content does not leave the phone) and adds where and when it was found and the takedown request; the whole's fingerprint is anchored to the public blockchain, with a complementary eIDAS seal (.asice). Three outcomes: the school file, the criminal complaint (Arts. 510 and 172 ter) and the platform's notice-and-action (DSA/UK OSA), with a sealed receipt. Limit: the seal does not remove the post, does not prevent the creation of hate and does not detect it; never a registry of the content or a deanonymization tool, only fingerprints on the device. Sources: the Picchio case; Molly Russell; ANAR; StopNCII.
- Figure 85.1 — Three shields for the elderly with a single logic (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, anchoring to the public blockchain, .asice): family verification against the cloned voice (challenge + sealed proof of life), attested care (anchored nursing home records; a camera only within the Supreme Court's proportionality doctrine) and the will with a certain date (Law 8/2021: support, not substitution). Limits: no central registry of voices, revocability, and control always in the hands of the elderly person. Sources: INCIBE; FBI/IC3; elDiario.es (El Escorial); Law 8/2021.
- Figure 86.1 — The issuing body (press office, operations center, central bank) computes on the device the SHA-512 fingerprint of the official statement — video, release, decision or alert — without the file leaving; for live statements by the principal it adds proof of life and forensic GPS. The whole's fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), and its verification handle (QR/URL) is published on the official channel. Two outcomes: on the left, public verification in seconds by journalists and citizens; on the right, the refutation of a forgery (hacked ticker, unverified rebroadcast) against the sealed original. Limit: the seal proves origin, time, author and integrity, not that the content is true, legal or sensible; it does not stop a deepfake never compared against it, nor protect the issuer from being hacked. Sources: the Ukraine 24/Zelensky case; the fake Pentagon image; the Rubio case; the Channel 14/Gallant case; the 2018 Hawaii alert; Article 50 of the AI Regulation; IPAWS; NSA/CISA Content Credentials CSI.
- Figure 87.1 — The genuine executive captures the payment instruction on the device (signed text, PDF invoice or video with proof of life and forensic GPS), computes its SHA-512 fingerprint without it leaving the device, and anchors it to the public blockchain, with a complementary RFC

3161/eIDAS timestamp (.asice). The treasury verifies the seal against the public anchor: if it checks out, it pays; if it is missing, it stops — the "urgent" instruction without a seal is, by design, a red flag (Arup's failure mode). Limit: the seal proves the act of authorizing, not that the instruction is legitimate; it stops neither a coerced executive nor a stolen key, and does not recover funds already transferred. Sources: the Arup case; FinCEN alert FIN-2024-Alert004; FBI/IC3 PSA250515; eIDAS Arts. 41-42.

- Figure 88.1 — Live identity with proof of existence (t2p: video selfie with random liveness challenges — blink, turn, smile, tilt — in <60 s, device biometrics or PIN, SHA-512 fingerprint, RFC 3161/eIDAS seal, anchor to the public blockchain, certificate with the video evidence and hardware attestation, .asice package verifiable offline): on the left, the meeting handshake (the meeting sheet as a challenge document signed by each participant before the microphones open; comparison of the sealed face with the on-screen one; the question against the current as an extra layer); on the right, signing up as a human (the sign-up form signed with a video selfie: pseudonymous-but-human, the platform verifies without holding biometrics; civil identity via the EUDI Wallet when needed). Context: real-time face-swap commercialized (Yahoo Boys), +2,665% native virtual camera attacks and +300% face swaps against liveness systems (iProov 2025), 30% projected enterprise distrust (Gartner). Limits: gesture challenges are a bar, not a wall — layers and asymmetric cost (the fraud that signs gets notarized); the face comparison belongs to the receiver, never the app. Sources: WIRED/9to5Mac; iProov; Gartner; FBI IC3; Decrypt; arXiv/MIT; EC (EUDI).
- Figure 89.1 — The issuer (or its agent or wire) seals the price-sensitive disclosure at the moment of issuing it: SHA-512 fingerprint on the device → public blockchain, with a complementary RFC 3161/eIDAS timestamp (.asice), with forensic GPS and proof of life if an executive is on camera. The fingerprint is published next to the official account; any journalist, exchange, regulator or algorithmic channel confirms in milliseconds whether a "breaking" tweet or release matches the genuine one (Lilly, Pentagon-Bloomberg, BSE, Ostin) or is the impostor. Limit: the seal proves who issued it and when, not that the disclosure is true; a sealed pump-and-dump release is still false, and manipulation remains the job of the SEC/ESMA/FINRA and the exchanges. Sources: the Eli Lilly case; the fake Pentagon image; the 2013 Hack Crash; the Ostin Technology case; the SEC/FINRA/NASAA alert; MAR Art. 12.
- Figure 90.1 — The beneficial owner, the director or the corporate service provider files the declaration and its evidence (shareholder register, trust deed, identity, proof of life with forensic GPS in the field); the SHA-512 fingerprint is computed on the device — the documents do not leave — and anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). The receipt complements the registries (FinCEN BOI · AMLR/AMLA · Companies House · FATF R.24): it proves who declared what and when, and makes backdating detectable. Limit: the seal does not prove that the declared person is the ultimate beneficial owner; it does not pierce chains of shell companies or detect a front man. Sources: Pandora Papers (ICIJ); Panama Papers; Transparency International; FinCEN; CJEU C-37/20 and C-601/20; FATF R.24; Open Ownership.
- Figure 91.1 — Remote hiring sealed in three links: video interview with proof of life → identity verification (document + selfie) → onboarding attestation bound to whoever receives the laptop. At each moment the SHA-512 fingerprint is computed on the device (the biometrics do not leave), with an RFC 3161/eIDAS seal, forensic GPS and anchoring to the public blockchain (.asice). The result: a non-repudiation chain proving that the same verified human was at all three points. Limit: the seal does not prove that the person is honest, that the identity is not stolen or that they are not a sanctioned national; it complements, does not replace, background and sanctions screening and deepfake detection. Sources: the KnowBe4 case; the Chapman case (DOJ); FBI/IC3 I-012325-PSA; NIST SP 800-63-4; Pindrop.
- Figure 92.1 — The platform employment relationship with proof of existence: the worker seals his work diary (hours, jobs, instructions, captures of conditions) on his device; the platform anchors the log of algorithmic decisions and the human review (Art. 22 GDPR; Directive (EU) 2024/2831); the works council receives the algorithm's parameters with sealed versioning (the Rider Law). Three outcomes: the presumption of employment proven with dated facts, the auditable deactivation and the inspection that does not start from zero. Limits: the seal does not decide employment status or replace collective bargaining; the diary is the worker's, never a reputational filter. Sources: OJEU (Dir. 2024/2831); BOE (Law 12/2021); elDiario.es (the Glovo case); Worker Info Exchange (Drivers v. Uber/Ola).
- Figure 93.1 — The working day with proof of existence (sealed with t2p — chapter 13): the worker's work diary (clock-ins, hours, shifts — on his device, only fingerprints on the chain), the clock-in with forensic GPS (the moment, never tracking), the roster and the order captured with a certain date,

and the company's official record certified in chained batches — provably intact before the Inspectorate and the court. Context: the digital time record reform (urgent processing since 30-9-2025) will demand immutability, traceability and remote access for the Inspectorate; >40% of overtime goes unpaid (EPA). Limits: the agreement and the judge decide the substance; proportionality always. Sources: EPA/INE; the digital time record draft.

- Figure 94.1 — Three irreversible milestones of a build: the pre-pour inspection (rebar, formwork, cover), the 28-day cube test, and the as-built drawings + geotechnical report at signing. At each milestone, t2p computes the SHA-512 fingerprint, adds forensic GPS and proof of life for field captures, seals with RFC 3161/eIDAS and anchors to the public blockchain (.asice); the container flows into the golden thread (BSA 2022) and the BIM model (ISO 19650) and pairs with the Digital Product Passport (Regulation EU 2024/3110). Risk points that independent dating makes provable: changing the computer's date (Hong Kong–Zhuhai–Macao) and writing the specification value (Testwell). Limit: the seal proves that a record existed, who made it and when, not that the building is safe or the material compliant. Sources: 2023 Turkey earthquakes (Isias Grand Hotel); NIST (Champlain Towers South); ICAC (HZMB bridge); Testwell; the Morandi bridge; Building Safety Act 2022; Regulation (EU) 2024/3110.
- Figure 95.1 — A maintenance event creates a record (EASA Form 1 / FAA 8130-3, logbook entry, odometer reading). The technician computes its SHA-512 fingerprint on the device (the file does not leave), adds forensic GPS and proof of life for field captures, and the whole's fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). Any later party — goods-in, CAMO, regulator, buyer — independently verifies who declared it and when. Limit: the seal proves the document's origin, time, author and integrity; it does not prove the part is genuine, new or airworthy, or the mileage true — garbage in, garbage on-chain. Sources: the AOG Technics case (SFO); CFM International's lawsuit; FAA estimates of unapproved parts; CARFAX and the European Parliament on odometers; the Aviation Supply Chain Integrity Coalition.
- Figure 96.1 — The chain of an emissions datum in four links: capture at the sensor (CEMS, methane analyzer, test bench, drone/satellite flight) → acquisition → regulatory report → audit. t2p computes on the device the SHA-512 fingerprint of the raw reading, adds an RFC 3161/eIDAS timestamp and, in the field, forensic GPS, and anchors it to the public blockchain (.asice); the auditor checks the submitted figure against the sealed capture. Two risk points: the rigged test bench (the VW/Cummins defeat device) and the sensor biased around the audits (Berkshire's CEMS), which the seal does NOT detect. Limit: the seal proves the measurement's origin, time, author and integrity; it does not prove the measurement true at the point of capture, quantify emissions or judge compliance. Sources: EPA (VW, Cummins, Berkshire, Marathon, the methane rule); IEA; MethaneSAT; CJEU (the thermal window).
- Figure 97.1 — Citizen science with proof of existence: the balcony sensor, the photo of the dumping and the noise diary are anchored at the origin (SHA-512 fingerprint, RFC 3161/eIDAS timestamp, public blockchain, .asice) and become chronologies of attested integrity that trigger the official inspection (a report with a homologated sound level meter, with a presumption of veracity), guide the expert and give memory to the legal voice of ecosystems (Mar Menor, Law 19/2022). Limits: the seal does not attest calibration or replace the official report; facts are documented, not people. Sources: CJEU (Doñana, 2021); BOE (Law 19/2022); Sensor.Community; the Ombudsman.
- Figure 98.1 — A battery's life cycle in four links: materials and mine → manufacture → use and state of health → second life / recycling. Each link generates an evidentiary artifact (recycling certificate, the notified body's carbon footprint report, the BMS export, the disassembly report); t2p computes its SHA-512 fingerprint on the device and anchors it to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), with forensic GPS and proof of life for field capture. The passport's QR links the claim to the sealed artifact, verifiable offline by customs, auditors and secondhand buyers. Limit: the seal proves the artifact's origin, date, author and integrity; it does not prove the recycled content real, the state of health untouched or the battery safe. Sources: Volvo EX90 / Circulor; the Global Battery Alliance; Regulation (EU) 2023/1542; the ESPR (EU) 2024/1781.
- Figure 99.1 — The waste chain with proof of existence (sealed with t2p — chapter 13): the container with place and time at every change of hands (forensic GPS; the sealing filmed with the date overlaid), the mass balance sealed in chained batches (dated inflows vs. outflows), the anchored final treatment certificate and the frozen green claim (web capture with source code). Context: Regulation (EU) 2024/1157, applicable from 21-5-2026 — prior notification for all exported plastic waste, a ban on non-OECD countries; Turkey received 39% of European exports (12.4 Mt). Limits: the trace attests to custody, not treatment quality; reduce and reuse first. Sources: OJEU/BOE; RETEMA.

- Figure 100.1 — The land with proof of existence (sealed with t2p — chapter 13): the logbook entry sealed in the moment and in the field (forensic GPS), the georeferenced photo the monitoring demands — with a neutral certain date — the season series as a defense against the satellite that errs (chapter 104's symmetry, seen from below) and the aid file in .asice. Context: the CAP watched by Sentinel/Copernicus; the digital logbook (CUE) voluntary since 2024, electronic plant protection records from 1-1-2027, the CUE mandatory on 1-1-2028 (>30 ha). Limits: it complements the official logbook, does not replace it; the countryside's digital divide demands cooperatives and agronomists as allies. Sources: FEGA/CUE rules; CAP 2023-2027.
- Figure 101.1 — The producer captures the geotagged field evidence with proof of life and forensic GPS (olive grove, hive, bottling line) and computes its SHA-512 fingerprint on the device; in parallel, the control body and the laboratory seal the certificate, the analytical report, the batch number and the PDO/PGI label. The bundle's fingerprint is anchored on the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), bound to a batch identifier. Any inspector, customs office or court verifies it offline: existence, author, date and integrity. Limit: the seal proves the paper and the field act, not that the oil is extra virgin or the honey uncut — the laboratory says that (NMR/IRMS/DNA/pollen). Sources: OPSON XIII and XIV (Europol); the Apulia case; Operation Omegabad; "From the Hives"; Reg. (EU) 2024/1143.
- Figure 102.1 — Two sources of evidence: the worker, the union or the field auditor captures on the device an image or a testimony of the workshop (SHA-512 fingerprint that does not leave the phone, forensic GPS, proof of life); the brand seals the supplier declaration and the due diligence record. Everything is sealed with an RFC 3161/eIDAS timestamp anchored on the public blockchain (.asice). Outcomes: worker testimony harder to dismiss as a deepfake, and a dated file for the CBP, the BAFA or the EU authority. Limit: the seal proves the act, the place and the date; it does not prove that the conditions are good, the declaration truthful or the testimony free. Sources: the Volkswagen/CBP case; the ILO; Miller & Chevalier (UFLPA); HRW (audits and automotive); the ECCHR (certificates).
- Figure 103.1 — When a fix is taken, an onboard unit or a port state inspector captures the navigation record (GNSS + time, IMO/MMSI, draft, AIS frame) and computes its SHA-512 fingerprint on the device; the whole's fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice), with forensic GPS and proof of life when there is a physical inspection. Two outcomes: on the left, a non-repudiable position-and-identity ledger; on the right, the absence of a seal becomes a red flag for OFAC/EU screening. Limit: the seal certifies that the device recorded this fix at this instant, not that the GNSS signal was genuine; a spoofed signal anchors a false position. Sources: the Windward case (the $800 million scheme); Kpler (261 vessels); GPSPATRON (10,000+ vessels); Scientific American (Hormuz); OFAC (MS ENOLA).
- Figure 104.1 — The Earth observation chain in two sealing links. At the origin, the operator (Maxar, Planet, Copernicus) computes the SHA-512 fingerprint of the raw/L1 scene plus metadata at ground station ingest; the fingerprint is anchored to the public blockchain, with a complementary RFC 3161/eIDAS seal (.asice). At the other end, the OSINT analyst or journalist seals again the crop, the annotation and the coordinates, with forensic GPS and proof of life. Risk points: imagery forged with GAN/diffusion (Stable Diffusion, StyleGAN2) and GNSS interference, which the Galileo OSNMA signature and the Vision Transformer detectors mitigate but do not eliminate. Limit: the seal proves origin, time, author and bit-for-bit integrity; it does not prove the scene authentic observation or the interpretation correct. Sources: the Bucha case (Maxar); Bo Zhao (deepfake geography); Xinjiang (BuzzFeed); arXiv 2511.17766; Galileo OSNMA; OGC IPT.
- Figure 105.1 — The emergency datum is sealed at the origin: the sensor reading (SAIH/Poyo), the coordinating body's minutes and video (CECOPI), the alert (ES-Alert), the evacuation order and the citizen's damage photo receive the SHA-512 fingerprint on the device or at the source, the RFC 3161/eIDAS timestamp and the anchoring on the public blockchain (.asice). Three outcomes: the judicial investigation (a non-repudiable chronology), the claim to the insurer and the Consortium, and public and journalistic auditing. Limit: the seal does not predict, does not warn and does not save; it only makes the truth indestructible. Sources: Maldita.es; CHJ/SAIH; the Insurance Compensation Consortium; Lauegi.
- Figure 106.1 — Against the often losing race of detection, the neutral anchor proves that something authentic existed and has not been altered.
- Figure 107.1 — The public AI decision leaves a cited record which, anchored with a proof of existence, becomes an independent, non-repudiable audit.
- Figure 108.1 — The personhood proof can be performed on the device, with no centralized biometric databases: one person, one voice, without surveillance.

- Figure 109.1 — The symmetrical proof in the street: the protester's video and the officer's body camera are sealed at capture, by segments (SHA-512 fingerprint on the device, RFC 3161/eIDAS timestamp, anchoring on the public blockchain, .asice), publishing nothing and identifying no one. In court, the two chains merge into a verifiable timeline; any cut or deletion is detected. Limits: the seal does not attest the context off-frame, does not weigh proportionality and demands revocability, voluntary publication and zero facial recognition. Sources: Pulitzer/Frazier; elDiario.es ("Rodea el Congreso"); SME/USPAC (Axon Body 3); Telemundo/Univision (Detroit).
- Figure 110.1 — The higher the temperature climbs, the further the economy falls: the cost of not acting soars with every tenth of a degree.
- Figure 111.1 — Each actor — government, company and citizens — has concrete actions for the short, medium and long term.
- Figure 112.1 — The seal and the chain: who controls the proof, and who gets trapped. The risks the chapter unpacks — the neutrality that does not exist, crystallized consent, the surplus value of the proof, the right to be forgotten deleted, unappealable algorithms and the self-censorship of the permanent record — and the counterweights it demands: hash-only by default, selective disclosure, democratic governance of the infrastructure and the right not to seal.
- Figure 113.1 — The proof under siege, from the lens to the courtroom (t2p: SHA-512 fingerprint on the device — the image never leaves the phone —, RFC 3161/eIDAS timestamp, anchor on the public blockchain, .asice package): the accredited photojournalist with an agency credential that does not expose the source; the citizen witness with commit-and-reveal (seal now, publish when it is safe) and posthumous proof; the three layers of the siege (door closed to the international press, local journalists killed or marked, "fake!" against real images) and where the seal breaks them. Limits: the blackout delays the anchor (the "not before / not after" temporal sandwich); spoofed GPS degrades the position, not the time; the seal proves when and what — never the truth of the scene: corroboration remains indispensable. Sources: CPJ; RSF; Costs of War (Brown); OHCHR; Forensic Architecture; ECCHR/Koblenz; Access Now; Starling Lab/Reuters; WITNESS/Guardian Project.
- Figure 114.1 — The Church's social doctrine facing the economy and technology: from 'Rerum Novarum' (1891, labor) to 'Laudato Si" (2015, the common home) and 'Magnifica Humanitas' (2026, the person in the age of AI).
- Figure 114.2 — The same proof of existence, two paths (§1): toward "Babel" (private power, the person reduced to data, domination) or toward the "city" (the dignity of those who exist, subsidiarity, universal access). The difference is not technical, but one of governance and purpose.

Glossary

Uso obligatorio en las tres versiones para garantizar coherencia. Las marcas **time2act** y **t2p no se traducen nunca**. Anglicismos asentados (blockchain, deepfake, smart contract, token) se mantienen.

Concepto	Català (maestro)	Español	English	Nota
Pico de demanda	pic de demanda	pico de demanda	peak demand	Sustituye al viejo "peak oil" (pico de oferta)
Pico de oferta	pic d'oferta (peak oil)	pico de oferta (peak oil)	peak oil	Tesis clásica de Roberts
Transición energética	transició energètica	transición energética	energy transition	
Prueba de existencia	prova d'existència	prueba de existencia	proof of existence	Núcleo de t2p
Sellado de tiempo	segellat de temps	sellado de tiempo	timestamping	
Resumen criptográfico	resum criptogràfic (hash)	hash criptográfico	cryptographic hash	"hash" admitido
Cadena de bloques	cadena de blocs / blockchain	blockchain	blockchain	"blockchain" de uso común
Libro mayor distribuido	llibre major distribuït	libro mayor distribuido	distributed ledger	
Contrato inteligente	contracte intel·ligent	contrato inteligente	smart contract	
Prueba de participación	prova de participació (PoS)	prueba de participación (PoS)	proof of stake (PoS)	Ethereum desde 2022
Prueba de trabajo	prova de treball (PoW)	prueba de trabajo (PoW)	proof of work (PoW)	Minería intensiva en energía
Procedencia del contenido	procedència del contingut	procedencia del contenido	content provenance	Estándar C2PA
Protocolo de confianza	protocol de confiança	protocolo de confianza	trust protocol	Marco de Tapscott
Coste normalizado de electricidad	cost normalitzat de l'electricitat	coste normalizado de la electricidad	LCOE	Siglas LCOE en las 3 lenguas
Red eléctrica	xarxa elèctrica	red eléctrica	(power) grid	
Red inteligente	xarxa intel·ligent	red inteligente	smart grid	
MRV de carbono	MRV de carboni	MRV de carbono	carbon MRV	Medición, reporte y verificación
Garantía de origen	garantia d'origen	garantía de origen	guarantee of origin	Certificados de energía renovable
Tokenización	tokenització	tokenización	tokenization	
Crédito de carbono	crèdit de carboni	crédito de carbono	carbon credit	1 crédito = 1 tCO_2e evitada o eliminada
Mercado de cumplimiento	mercat de compliment (EU ETS)	mercado de cumplimiento (EU ETS)	compliance market (EU ETS)	Cap-and-trade regulado; el derecho es la EUA
CBAM	mecanisme d'ajust de carboni a la frontera (CBAM)	mecanismo de ajuste de carbono en frontera (CBAM)	Carbon Border Adjustment Mechanism (CBAM)	Fase definitiva desde el 1-ene-2026
CORSIA	CORSIA	CORSIA	CORSIA	Compensación de emisiones de la aviación (OACI)

Concepto	Català (maestro)	Español	English	Nota
Mercado voluntario	mercat voluntari del carboni (VCM)	mercado voluntario de carbono (VCM)	voluntary carbon market (VCM)	Verra, Gold Standard; sin obligación legal
Doble cómputo	doble còmput	doble cómputo	double counting	Contar el mismo ahorro o crédito dos veces
Crédito fantasma	crèdit fantasma	crédito fantasma	phantom credit	Crédito sin una reducción real detrás
Adicionalidad	addicionalitat	adicionalidad	additionality	¿Habría ocurrido igualmente sin el incentivo?
Línea de base	línia de base	línea de base	baseline	Escenario de referencia para medir el ahorro
Principios de Carbono de Base	ICVCM / Core Carbon Principles	ICVCM / Core Carbon Principles	ICVCM / Core Carbon Principles	Sello de integridad del mercado voluntario (2023)
Certificado de Ahorro Energético	CAE (Certificat d'Estalvi Energètic)	CAE (Certificado de Ahorro Energético)	Energy Saving Certificate (white certificate)	1 CAE ≈ 1 kWh acreditado; Sistema CAE (RD 36/2023)
Sujeto obligado	subjecte obligat / delegat	sujeto obligado / delegado	obligated / delegated party	Comercializadoras energéticas y sus delegados
Telemática	telemàtica	telemática	telematics	Datos GPS/CAN-bus del vehículo
Coche compartido	cotxe compartit	coche compartido	car-sharing	Un vehículo compartido sustituye a varios privados
tCO_2e	tCO_2e	tCO_2e	tCO_2e	Tonelada de CO_2 equivalente
Activos varados	actius varats	activos varados	stranded assets	Reservas fósiles no quemables
Hidrógeno verde	hidrogen verd	hidrógeno verde	green hydrogen	
Captura de carbono	captura de carboni (CCS)	captura de carbono (CCS)	carbon capture (CCS)	
Cero neto	zero net	cero neto	net zero	
Gases de efecto invernadero	gasos d'efecte hivernacle (GEH)	gases de efecto invernadero (GEI)	greenhouse gases (GHG)	
Vehículo eléctrico	vehicle elèctric (VE)	vehículo eléctrico (VE)	electric vehicle (EV)	
Ultrafalsificación	deepfake	deepfake	deepfake	Anglicismo; glosar 1ª vez
Desinformación	desinformació	desinformación	disinformation	
Credenciales de contenido	credencials de contingut (C2PA)	credenciales de contenido (C2PA)	Content Credentials (C2PA)	Estándar de procedencia firmada
Sello de tiempo RFC 3161	segell de temps RFC 3161	sello de tiempo RFC 3161	RFC 3161 timestamp	Sellado de tiempo de confianza
OpenTimestamps	OpenTimestamps	OpenTimestamps	OpenTimestamps	Sellado en Bitcoin vía árbol de Merkle
Identidad digital europea	eIDAS 2.0	eIDAS 2.0	eIDAS 2.0	Reglamento UE de identidad y confianza
Cartera de identidad digital europea	cartera EUDI	cartera EUDI	EUDI Wallet	Monedero europeo de identidad (eIDAS 2.0); cada estado debe

Concepto	Català (maestro)	Español	English	Nota
				ofrecerla desde dic-2026
Conozca a su cliente	KYC	KYC	KYC (know your customer)	Verificación de identidad en finanzas
Credencial verificable	credencial verificable	credencial verificable	verifiable credential	Modelo W3C; diplomas y títulos
Cadena de custodia	cadena de custòdia	cadena de custodia	chain of custody	Integridad probatoria
Protocolo de Berkeley	Protocol de Berkeley	Protocolo de Berkeley	Berkeley Protocol	Estándar OHCHR/Berkeley para evidencia digital de fuentes abiertas
Compromiso y revelación	compromís i revelació (commit-reveal)	compromiso y revelación (commit-reveal)	commit-reveal	Sellar hoy la huella, revelar el contenido cuando convenga
Investigación de fuentes abiertas	recerca en fonts obertes (OSINT)	investigación de fuentes abiertas (OSINT)	open-source intelligence (OSINT)	Verificación con material público: geolocalización, cronolocalización
Ataque de inyección	atac d'injecció	ataque de inyección	injection attack	Cámara virtual que burla el KYC
Dividendo del mentiroso	dividend del mentider	dividendo del mentiroso	liar's dividend	Eludir la responsabilidad negando como deepfake una grabación auténtica
Árbol de Merkle	arbre de Merkle	árbol de Merkle	Merkle tree	Raíz que resume miles de hashes en uno
Función SHA-512	SHA-512	SHA-512	SHA-512	Huella de 512 bits que usa t2p
Polygon	Polygon	Polygon	Polygon	Cadena pública (compatible con Ethereum) donde t2p ancla hoy; la implementación admite cualquier otra L2
Contenedor ASiC-E	contenidor ASiC-E (.asice)	contenedor ASiC-E (.asice)	ASiC-E container (.asice)	Paquete de evidencia verificable offline (ETSI EN 319 162)
Prueba de vida	prova de vida	prueba de vida	liveness / proof of life	Biometría anti-suplantación (vídeo-selfie)
Hash perceptual	hash perceptual	hash perceptual	perceptual hash	Huella robusta a recompresión o reescalado
Fraude de siniestros	frau de sinistres	fraude de siniestros	claims fraud	Reclamaciones de seguro infladas o inventadas
Falsificación	falsificació	falsificación	counterfeiting	Bienes o documentos imitados ilegalmente
Trazabilidad de origen	traçabilitat d'origen	trazabilidad de origen	provenance / traceability	Prueba verificable de dónde viene un bien
EUDR	EUDR (reglament antidesforestació)	EUDR (reglamento antideforestación)	EU Deforestation Regulation	Exige geolocalización de las parcelas de origen

Concepto	Català (maestro)	Español	English	Nota
Registro de la propiedad	registre de la propietat	registro de la propiedad	land registry	Padrón oficial de títulos de tierra
Activo del mundo real	actiu del món real (RWA)	activo del mundo real (RWA)	real-world asset (RWA)	Activo físico/financiero representado on-chain
Problema del oráculo	problema de l'oracle	problema del oráculo	oracle problem	Cómo probar que el dato off-chain es cierto
Publicación defensiva	publicació defensiva	publicación defensiva	defensive publication	Difundir una invención para crear prior art
Prior art / estado de la técnica	estat de la tècnica	estado de la técnica	prior art	Conocimiento previo que impide patentar
Contratación pública	contractació pública	contratación pública	public procurement	Compra de bienes y obras por el Estado
Manipulación de licitaciones	manipulació de licitacions	manipulación de licitaciones	bid-rigging	Acuerdo fraudulento para amañar un concurso
Identidad autosoberana	identitat autosobirana (SSI)	identidad autosoberana (SSI)	self-sovereign identity (SSI)	El usuario custodia sus propias credenciales
Identificador descentralizado	identificador descentralitzat (DID)	identificador descentralizado (DID)	decentralized identifier (DID)	Estándar W3C de identidad sin autoridad central
Apátrida	apàtrida	apátrida	stateless person	Persona sin nacionalidad reconocida por ningún Estado
Auditoría de riesgo limitado	auditoria de risc limitat (RLA)	auditoría de riesgo limitado (RLA)	risk-limiting audit (RLA)	Recuento estadístico que confirma el resultado electoral
Asistencia en efectivo y vales	assistència en efectiu i vals (CVA)	asistencia en efectivo y vales (CVA)	cash & voucher assistance (CVA)	Ayuda humanitaria entregada como dinero o vales
Deduplicación de beneficiarios	deduplicació de beneficiaris	deduplicación de beneficiarios	beneficiary deduplication	Evitar el doble registro de un mismo beneficiario
Testamento ológrafo	testament hològraf	testamento ológrafo	holographic will	Testamento escrito a mano por el testador
Conocimiento de embarque	coneixement d'embarcament	conocimiento de embarque	bill of lading (B/L)	Título de la mercancía transportada
Carta de crédito	carta de crèdit	carta de crédito	letter of credit (L/C)	Garantía bancaria de pago en el comercio internacional
Documento electrónico transferible	document electrònic transferible (MLETR)	documento electrónico transferible (MLETR)	electronic transferable record (MLETR)	Ley modelo UNCITRAL que da validez legal a los documentos electrónicos
Aseguramiento	assegurament	aseguramiento	assurance	Verificación independiente de un informe (financiero o de sostenibilidad)
CSRD	CSRD	CSRD	CSRD	Directiva UE de informes de sostenibilidad con aseguramiento obligatorio

Concepto	Català (maestro)	Español	English	Nota
Sitio hecho para anunciar	lloc fet per anunciar (MFA)	sitio hecho para anunciar (MFA)	made-for-advertising (MFA)	Web basura creada solo para captar gasto publicitario
Tiempo de permanencia	temps de permanència	tiempo de permanencia	dwell time	Días que un atacante pasa dentro de una red sin ser detectado
Registro a prueba de manipulación	registre inalterable	registro inalterable	tamper-evident log	Log cuyo hash anclado delata cualquier alteración
Procedencia de la obra	procedència de l'obra	procedencia de la obra	art provenance	Cadena verificable de propiedad y origen de una obra
Divulgación selectiva	divulgació selectiva	divulgación selectiva	selective disclosure	Probar un dato sin revelar el resto (zero-knowledge)
Integridad de la investigación	integritat de la recerca	integridad de la investigación	research integrity	Integridad del dato y la imagen científicos, no de la autoría
Pre-registro de ensayos	preregistre d'assaigs	preregistro de ensayos	trial pre-registration	Depósito con fecha cierta del protocolo antes de recoger datos
Fábrica de papers	fàbrica de papers	fábrica de papers	paper mill	Empresa que produce y vende artículos científicos falsos
Prueba de vida recurrente	prova de vida recurrent	prueba de vida recurrente	recurring proof of life	Atestación periódica de que el beneficiario sigue vivo
Pasaporte biológico	passaport biològic	pasaporte biológico	biological passport	Serie longitudinal de biomarcadores del atleta (detección indirecta de dopaje)
Amaño de partidos	amanyament de partits	amaño de partidos	match-fixing	Manipulación del resultado deportivo, detectada por patrones de apuestas
Fraude de streaming	frau de reproduccions	fraude de streaming	streaming fraud	Bots que inflan reproducciones para drenar el fondo de royalties
Intercambio de SIM	intercanvi de SIM	intercambio de SIM	SIM swapping	Robo del número de móvil para interceptar códigos de un solo uso
Autenticación del identificador	autenticació de l'identificador	autenticación del identificador	STIR/SHAKEN	Firma del número de origen de una llamada en la red
Componente reciclado/remarcado	component reciclat/remarcat	componente reciclado/remarcado	recycled/remarked component	Chip usado, lijado y revendido como nuevo
Pedigrí del componente	pedigrí del component	pedigrí del componente	component pedigree	Cadena de custodia documentada hasta el fabricante original
Pesca INDNR	pesca INDNR	pesca INDNR	IUU fishing	Ilegal, no declarada y no reglamentada
Minerales de conflicto (3TG)	minerals de conflicte (3TG)	minerales de conflicto (3TG)	3TG conflict minerals	Estaño, tántalo, wolframio y oro

Concepto	Català (maestro)	Español	English	Nota
Proceso de Kimberley	Procés de Kimberley	Proceso de Kimberley	Kimberley Process	Certificación contra los diamantes de sangre
Transparencia de donaciones	transparència de donacions	transparencia de donaciones	donation transparency	Rastreo del dinero del donante hasta el impacto demostrable
Registrador de datos (caja negra)	registrador de dades (caixa negra)	registrador de datos (caja negra)	data recorder (black box)	FDR/CVR/EDR/log de vuelo a prueba de manipulación
Dato genómico	dada genòmica	dato genómico	genomic data	Categoría especial; irrevocable y compartido con la familia
Recibo de consentimiento	rebut de consentiment	recibo de consentimiento	consent receipt	Constancia legible y procesable de un acto de consentimiento (Kantara/ISO 27560)
Registro de auditoría	registre d'auditoria	registro de auditoría	audit trail	Rastro de quién accedió o alteró un registro electrónico y cuándo
Lista de materiales del modelo	llista de materials del model (AI-BOM)	lista de materiales del modelo (AI-BOM)	AI/ML bill of materials (AI-BOM)	Inventario de datos, pesos y evaluaciones de un modelo (CycloneDX ML-BOM)
Tarjeta de modelo	targeta de model	tarjeta de modelo	model card	Documento de uso, rendimiento y límites de un modelo (Mitchell 2019)
Réplica digital	rèplica digital	réplica digital	digital replica	Recreación por IA de la voz o la imagen de una persona
Derecho a la propia imagen	dret a la pròpia imatge	derecho a la propia imagen	right of publicity / likeness rights	Control del uso de cara y voz (ELVIS Act, NO FAKES, ley danesa)
Imágenes íntimas no consentidas	imatges íntimes no consentides (NCII)	imágenes íntimas no consentidas (NCII)	non-consensual intimate imagery (NCII)	Difusión de imágenes sexuales sin consentimiento, reales o deepfake (TAKE IT DOWN Act)
Huella en el dispositivo	empremta al dispositiu	huella en el dispositivo	on-device hashing	El archivo no sale del dispositivo; solo viaja la huella (modelo StopNCII)
Fraude del directivo	frau del directiu (BEC)	fraude del directivo (BEC)	business email compromise (BEC)	Suplantación de un alto cargo para autorizar un pago (caso Arup)
Verificación fuera de banda	verificació fora de banda	verificación fuera de banda	out-of-band verification	Confirmar una orden por un canal independiente y pactado
Titular real último	titular real últim (UBO)	titular real último (UBO)	ultimate beneficial owner (UBO)	Persona física que controla de verdad una sociedad (GAFI R.24)
Sociedad pantalla	societat pantalla	sociedad pantalla	shell company	Empresa sin actividad real usada para ocultar al titular

Concepto	Català (maestro)	Español	English	Nota
Hilo dorado	fil daurat	hilo dorado	golden thread	Registro digital trazable del edificio (Building Safety Act 2022)
Planos as-built	plànols as-built	planos as-built	as-built drawings	Registro de cómo se construyó realmente, no de cómo se diseñó
Dispositivo de desactivación	dispositiu de desactivació	dispositivo de desactivación	defeat device	Trampa que falsea las emisiones durante la prueba (Dieselgate)
Monitorización continua de emisiones	monitoratge continu d'emissions (CEMS)	monitorización continua de emisiones (CEMS)	continuous emissions monitoring (CEMS)	Sensores que reportan emisiones en tiempo real
Pasaporte Digital de Producto	Passaport Digital de Producte (DPP)	Pasaporte Digital de Producto (DPP)	Digital Product Passport (DPP)	Registro de ciclo de vida (ESPR UE; batería obligatoria desde 2027)
Pieza sospechosa no aprobada	peça sospitosa no aprovada (SUP)	pieza sospechosa no aprobada (SUP)	suspected unapproved part (SUP)	Componente con certificado de aeronavegabilidad falsificado (AOG Technics)
Fraude alimentario	frau alimentari	fraude alimentario	food fraud	Adulteración o falsa indicación de origen (Operación OPSON)
Indicación geográfica	denominació d'origen (DOP/IGP)	denominación de origen (DOP/IGP)	geographical indication (PDO/PGI)	Sello de origen y método protegidos en la UE
Diligencia debida en la cadena	diligència deguda (CSDDD)	diligencia debida (CSDDD)	supply-chain due diligence (CSDDD)	Deber de vigilar el trabajo forzoso y los derechos humanos (UFLPA, CSDDD)
Falseamiento del AIS	falsejament de l'AIS	falseamiento del AIS	AIS spoofing	Falsear la posición o la identidad de un buque
Flota fantasma	flota fantasma	flota fantasma	shadow / dark fleet	Buques que evaden sanciones ocultando posición e identidad
Geografía deepfake	geografia deepfake	geografía deepfake	deepfake geography	Imagen de satélite sintética o manipulada; la imagen orbital como testigo y como blanco
Handshake de vida	handshake de vida	handshake de vida	liveness handshake	Intercambio de vídeo-selfies sellados sobre un documento-reto antes de una reunión o alta (cap. 88)
Credencial de personalidad	credencial de personalitat	credencial de personalidad	personhood credential	Probar que eres humano real sin revelar identidad (propuesta MIT/OpenAI/Microsoft 2024)