

The Other Bill

The material footprint of AI, the collapse of the cost of reading, and the decade in which both arrive

Jan Knacksterdt

August 2026

Abstract

The build-out of AI computing infrastructure is driving one of the largest material procurements of the decade: predominantly copper, aluminium and the electrical infrastructure behind each gigawatt, alongside smaller tonnages with sharper chains in heavy rare earths, battery metals and gallium. These materials originate in supply chains with well-documented and structurally stable problems, from water conflict and tailings risk to the near-total absence of oversight in artisanal mining, conflict-financed rare earth extraction and midstream processing. This essay argues that the same industry is simultaneously producing the instrument that ends the era of the unexamined supply-chain claim: the marginal cost of document analysis is falling roughly tenfold per year toward zero, converting corporate accounts written for skimming readers into liabilities awaiting retroactive examination. By default, this capability flows to the paying side of every dispute. The essay develops the case that AI companies carry a specific, non-philanthropic responsibility for how it is routed, toward affected communities, Indigenous organisations, workers and the civil-society organisations that hold the truth on the ground, and sets out five components of that responsibility, illustrated by what existing accountability architectures such as IRMA and worker-driven monitoring have been waiting for.

Something is about to happen to every company with a supply chain, and the AI industry is doing it to them. For thirty years, a sustainability claim was safe by default, because checking it against the ground was expensive. That protection is ending: the marginal cost of document analysis is going to zero, and with it the era in which a company could say things about its supply chain that nobody would ever line-check. Every company will feel this, but one industry sits on both sides of it at once, driving one of the largest material build-outs of the decade into some of the world's most difficult supply chains, while building the very instrument that brings that era to a close. This essay is about that industry, that bill, and the decision it has not yet noticed it is making.

I. The Bill

Start with the numbers everyone knows, and then keep going one step further than the conversation usually does.

Global datacentre electricity consumption stood at roughly 415 TWh in 2024. The IEA's base case puts it at about 945 TWh by 2030, more than doubling in six years, growing four times faster than every other sector of electricity demand combined.¹ Installed datacentre capacity has reached roughly 100 gigawatts; McKinsey projects AI workloads alone requiring around 156 GW by 2030.² Individual campuses are being announced at 5 GW, a single site consuming in a year around three times what the city of New Orleans consumes.³

Everyone in the industry knows these numbers. They are the numbers of the scaling debate: can the grid keep up, will the power arrive in time, who gets the interconnects. The IEA itself warns that a fifth of planned projects face delay because transmission cannot be built fast enough.⁴

There is a step the conversation doesn't take, which is to ask what a gigawatt is actually made of.

A gigawatt of datacentre is not primarily chips. It is buildings, switchgear, transformers, busbars, cabling, cooling plant. And beyond the fence, everything the grid must add to deliver that gigawatt: transmission lines, distribution, substations, generation, storage. Copper intensity inside the facility runs at roughly 27 to 33 tonnes per megawatt of applied power.⁵ Industry estimates put around 700,000 tonnes of copper into datacentre construction through 2030, and as much as five million tonnes into the transmission and distribution build-out that powers it, of which roughly 1.1 million tonnes is directly attributable to datacentres.⁶ The material bill of the AI build-out is, by a factor of several, an *electrical infrastructure* bill.

Copper is the anchor, but it does not travel alone. Aluminium follows it through the transmission lines and the structures. Steel and concrete carry the buildings and the turbine towers. And then come the smaller tonnages with the sharper chains. The permanent magnets in cooling systems, in backup generators, in the hard drives that hyperscale storage still runs on, and above all in the wind turbines increasingly contracted to power the campuses. These are neodymium-iron-boron magnets, and the heavy rare earths that let them survive heat, dysprosium and terbium, come today principally from one place, which Part II will visit. The batteries in the uninterruptible power supplies and the grid-scale storage co-located with the campuses run on lithium, nickel and cobalt. The power electronics increasingly run on gallium. The tonnages are small next to copper. The chains are not.¹⁹

Every gigawatt announced on a stage is a procurement order into this system.

There is a standard for whose responsibility that is, and it is not obscure. The UN Guiding Principles on Business and Human Rights, endorsed unanimously in 2011, cover exactly this case: a company connected to harms through its business relationships, without causing them. The standard does not ask for control of the supply chain. It asks for something more uncomfortable: that you use the influence

you have, *build* the influence you lack, and say publicly where it runs out.⁷ “We don’t own the mines” is not an exemption under the framework. It is the opening sentence of the obligation.

One distinction matters here, because “the AI industry” covers two positions. Those who build and own the infrastructure hold the purchasing relationships, specify the equipment and choose the sites, and most of them already run responsible sourcing functions. Those who buy compute trigger the procurement without carrying it out. At the largest firms both positions sit in the same house. The obligations differ in what they demand, not in whether they apply, and Part VI comes back to that.

A qualification, so it is on the record: datacentres remain a smaller copper driver than electrification, grids and electric vehicles. That limits the scale of the claim. It does not work as a justification, because “our share is small” is the argument every contributing sector makes simultaneously, and accepting it would exempt them all.

So the situation, stated plainly: the industry is contracting for one of the largest material build-outs of the decade, into supply chains whose problems Part II will lay out, under a responsibility standard its own scale activates, and its sustainability conversation is about PUE and power purchase agreements.

So much for the bill. Before getting to the strange part, it is worth looking at where these materials come from.

II. The Chains

The problems in these chains are not scandals awaiting exposure. They are structures, documented for years, sometimes decades, and stable precisely because engaging with them has been expensive. Six of them carry the weight, and the last two are the ones the corporate conversation most reliably skips.

Water and the living world

The best copper sits under the driest ground. In the Atacama and the Peruvian highlands, mines and communities draw on the same aquifers, and the accounting of who draws what is contested at every step: abstraction rights on paper, actual drawdown in the ground, and community observation of wells and wetlands that enters no official record. Water disputes are among the most common causes of suspended mining operations worldwide.

There is another face to the problem, one that files no complaint at all. Further along the battery chain, much of the nickel now being dug in Indonesia comes out from under some of the most species-rich rainforest on earth; on Halmahera, geospatial analysis has identified at least 5,331 hectares of tropical forest cleared inside nickel mining concessions, in a landscape that is also home to one of the last peoples living in voluntary isolation anywhere.¹³ None of this is hidden. That figure exists because somebody sat down with satellite imagery and concession boundaries and did the work.

The imagery had been there all along, updated every few days, in archives anyone can access. What has been missing is not the data but the reading: someone to join the deforestation polygon to the concession that caused it, the water series to the permit that allowed it, the species record to the licence condition that promised otherwise.

Nature generates evidence continuously and pays no one to process it.

And the pattern is not confined to the mines: the datacentres at the other end of the chain repeat the water half of it. A fifth of US servers' direct water footprint already comes from moderately to highly stressed watersheds, and nearly half of servers draw power from plants in water-stressed regions.⁸ It is the same asymmetry, at both ends of the same chain.

Tailings

For every tonne of copper, hundreds of tonnes of finely ground rock, stored wet behind earthen embankments that must stand forever. Tens of thousands of these facilities exist; a meaningful fraction are aging, inactive, or orphaned: no solvent owner, no monitoring, no one responsible. After a dam failure at Brumadinho killed 272 people in January 2019, investors forced the creation of a global disclosure database; roughly 1,800 facilities are in it, on the basis of *voluntary* reporting by the companies asked.⁹ The industry standard that followed covers the organised large-operator segment. The facilities most likely to fail are, by construction, in neither.

Land, and the people already on it

This is not a marginal overlap. A study in *Nature Sustainability* geolocated 5,097 mining projects across the roughly thirty minerals of the energy transition, the same minerals in Part I's bill, and found that 54 percent sit on or near the lands of Indigenous peoples, with more than half of the total resource base on or near Indigenous or peasant land. For lithium, 85 percent of reserves and resources overlap Indigenous lands; for manganese, three quarters.¹⁰

These are the two groups whose rights to consultation and to free, prior and informed consent are written into United Nations declarations. The architecture for exercising those rights exists everywhere on paper: impact assessments, comment periods, consent processes. In practice, consent runs through documents (four hundred pages, technical register, second language, sixty days), and "informed" is a condition that is asserted rather than established. The proceedings are formally open and practically closed, which is how a valley loses its water without a single recorded objection. Hold that mechanism in mind; Part III is about what just happened to it.

The forty-five million

Artisanal and small-scale mining needs its own entry, because it is neither a labour problem nor a sourcing problem but an economy. Roughly 45 million people dig directly, for cobalt, gold, tin, tantalum and tungsten among others, across some eighty countries; counting dependents and downstream trades, an estimated 315 million

people live from it, on figures collected around 2019 and 2020.¹¹ That is more than the population of Indonesia, living on the informal edge of the same chains that feed the formal bill in Part I.

The defining condition is informality, and it is worth spelling out what that means in practice. Most of this mining happens without licences, on land the miners hold no title to, selling into cash chains through aggregators who set the price at the pit's edge. No licence means no legal existence: no access to formal buyers, no bank account for the cooperative that doesn't exist, no standing when a concession holder or an armed unit moves in, no safety rule that anyone enforces. The work is dangerous in the ways documented for twenty years, from tunnel collapses to mercury exposure, and the people doing it capture a fraction of the value of what they dig.

Formal supply chains have mostly responded by pushing the sector away. Clean-chain policies that accept certified industrial mines only sound responsible and do the opposite: the material still arrives, laundered through aggregators, while the miners lose the one lever they had, which is access to buyers who ask questions. The framework itself warns against this reflex. Withdrawal that worsens harm is not responsibility, and a company that stays must be able to show what it is doing with the influence it kept.⁷ What engagement instead of exclusion looks like, and why it has never managed to scale, is a Part V question. Here the point is only the size of the gap: the largest workforce anywhere in this bill is the one with the least access to any of its formal structures.

The chain nobody can audit

The rare earths deserve their own entry, because they are the corner of the bill where every problem in this essay compounds at once.

The heavy rare earths in heat-resistant magnets, dysprosium and terbium, come today principally from northern Myanmar. Since China tightened its own environmental rules, extraction migrated across the border into Kachin State: close to 400 mining sites and nearly 3,000 leaching ponds, operated in territory controlled by armed groups, in a country whose formal economy collapsed after the 2021 coup. The method is in-situ leaching: ammonium sulphate injected through pipe networks into forested mountainsides, the solution collected in lined ponds, the sludge fired in wood-fuelled furnaces. Workers stir acid pools protected by raincoats. Landslides follow the deforestation. No environmental impact assessment has ever been conducted for any rare earth operation in the country. Between 2017 and 2024, more than 290,000 tonnes moved into China, roughly two thirds of China's rare earth imports, and from there, through separation and magnet-making, into wind turbines, vehicles, drives and cooling systems worldwide.¹²

Note what this case does to the standard corporate toolkit. There is no operator to engage, no audit to commission, no certification to require. The territory is not accessible to any of it. Here, the only forms of due diligence that exist at all are the remote ones: satellite observation, customs data, tracing work by organisations like the ones that documented everything just cited. For this corner of the bill, in other words,

the remote methods of Part III are not an improvement on the current system. There is no other system.

And the harm does not stop at the border. Further south, along the rivers that run out of Shan State into northern Thailand, the same extraction has become a downstream problem with a paper trail. A toxicologist at Chiang Mai University sampled seven points along the Kok river and found arsenic above the standard at all of them, rising toward the Myanmar border, alongside cadmium, chromium and mercury; Thailand's own pollution authority ran nine monitoring rounds to August 2025, and the first eight exceeded the arsenic limit at every checkpoint. Satellite analysis places rare earth mines immediately across the border. An independent Thai newsroom put the cost to twenty communities in Chiang Rai at 511 million baht in agriculture and direct income, and around 1.3 billion baht a year once fishing and tourism are counted.¹² Fishermen have stopped fishing. One farmer stripped the topsoil off his fields and still lost half his income.

The midstream blind spot

Now the gap the reporting shelf reveals by what is not on it. The public conversation about mineral chains is a conversation about *mining*. Between the pit and the product sits the midstream: smelting, refining, separation, precursor chemistry. It is at once the most concentrated link in every chain and the least examined.

The concentration is strategic fact: the large majority of the world's cobalt refining, rare earth separation and battery-grade processing runs through a small number of industrial clusters, most of them in China or Chinese-financed. The scrutiny gap is just as much a fact. Mining has decades of NGO documentation, community organisation and litigation attached to it. Processing has fences, industrial parks, and workforces that are easier to reach with a contract than with a survey.

Where independent reporting has reached the midstream, what it found justifies the worry. The Indonesia Morowali Industrial Park, the world's largest nickel processing complex feeding stainless steel and battery chains, recorded, across Indonesian nickel facilities, 114 accidents with 101 deaths and 240 injuries between 2015 and mid-2024, by a civil-society count assembled from media reports because no official register exists. A furnace explosion in December 2023 killed 21 workers; investigations found retaliation against workers who raised safety concerns; more than two dozen workers across nine companies described a production-first culture to journalists, backed by three unions representing twelve thousand members.¹³ The workers organised anyway; a union declared itself two months after the explosion.

The structural point is bigger than one park. Midstream harms are *labour* harms: hours, safety, retaliation, subcontracting. The audit machinery of the sourcing world is systematically bad at them, because it interviews workers on site, in working hours, within earshot of the employment relationship. An announced audit measures a facility's ability to pass an announced audit. The people who know the truth about a smelter are the eighty thousand people inside it, and no assurance model in current use gives their knowledge a protected route into the record. Worker-driven models

exist and have a track record in other sectors; the midstream of the AI bill has almost none of them. This is the least-lit link of the chain, and Part IV will come back to why that specific darkness is now a choice.

The account and the ground

The sixth problem contains the other five, and it has three distinct layers that are usually blurred together.

The first layer is *corporate self-report*. Every company in these chains publishes: sustainability reports, tailings disclosures, water accounts, human rights statements. Many of these documents are careful, and the better ones are considerably more candid than they are given credit for. What none of them can do is close the distance to the specific facility, the specific aquifer, the specific comment period, because the evidence that would close it was priced out of reach.

That limit is structural, not a failure of the people who write them. Ask anyone who has produced a supply chain report and they will tell you precisely which passages rest on a supplier's own declaration, which country risk section could not be checked against local sources, and which site they would have visited with another two weeks and another budget line. The authors are usually the best-informed critics of their own documents. The reports look the way they do because verification cost what it cost, and because the reader they were written for, a rating agency working through indicators, was the only reader who reliably showed up. Behind some of these documents sits real work: teams that have spent years at a single site, engagement that outlived several reporting cycles, projects that left the annual rhythm behind entirely. The claim here is not that nothing happens. It is that the format cannot show which part of it holds, and never could. Part V returns to that work and to why there is so little of it.

The second layer was built to fix the first, and it is where the finding gets uncomfortable. Certification schemes, industry standards, multi-stakeholder initiatives: the *assurance industry* that grew up around corporate self-report over three decades. A ten-year study of forty such initiatives, covering more than ten thousand corporate members, concluded that the experiment failed at its protective purpose: fewer than one in five include affected people in their governance at all, and only five of the twenty oldest had ever measured their effect on the people they exist to protect.¹⁴ The audits are periodic, announced, and paid for by the audited. And even the best-built mechanisms show a structural signature: researchers examining one of the field's most developed complaint systems found it scored well on *every* procedural criterion while producing poor outcomes in the actual cases.¹⁵ Process conformity and delivery had come apart, and the assessment framework could not see it.

A third layer has grown up more recently: public comparative scorecards such as Lead the Charge or KnowTheChain, iterated annually by coalitions of NGOs and investors, which rank companies against each other and even grade the certification schemes they rely on. They are a real advance in comparability, and they rest, by their own methodology, on published company reporting. Which is to say: even the newest accountability method reads the account, not the ground, because reading the ground

has been unaffordable.¹⁴

So the machinery of accountability exists, three layers deep now. One produces the account, one certifies it, one compares it, and all three stop at the same place, because the step beyond it was never affordable. Each layer was built to work around that limit rather than to remove it. The interesting question is what the machinery does once the limit is gone. This is not a case for tearing the machinery down; thirty years of institution-building sit in it, and Part V will visit the corner of it that got the architecture right. It is a case for seeing precisely what the machinery has always lacked: continuous evidence, protected channels for the people who hold the truth on the ground, and reading capacity on the affected side. All three have just changed in price.

Six problems. Notice what every one has as its *stabiliser*: engagement is expensive. Reading the assessment, monitoring the dam, checking the report against the satellite record, reaching the workers inside the fence, staying on the case: all of it priced in specialist attention, and the price has protected the equilibrium for thirty years.

Which is why what follows changes everything in this part, and not only for the AI industry.

III. The Collapse


While the material curve compounds upward, a second curve is collapsing downward, and almost nobody has connected the two.

For thirty years, sustainability claims have lived inside a cost asymmetry. Checking a company's account of itself against permits, court records, satellite imagery and local press in thirty languages required a research team and months. The number of organisations that could do it was small, their attention was scarce, and every company knew it. The result was the equilibrium Part II ended on: generous commitments at the level of principle, confident reporting at the level of the PDF, and a vanishingly small probability that anyone would ever line the two up against the facts on the ground.

That asymmetry is not eroding. It is collapsing, at a rate with no precedent in any prior technology transition.

The cost of machine reading, which is what verification is made of, is falling roughly tenfold per year at constant capability. GPT-4 launched in March 2023 at \$30 per million input tokens and \$60 per million output; equivalent capability, from whichever provider offers it most cheaply, now runs below fifty cents. One capability tier down, the price fell roughly a thousandfold in three years. Benchmark-level analysis finds prices for fixed performance milestones falling between 9x and 900x per year, median around 50x: the cost of a fixed quality of analysis halving every two months.¹⁷ That is faster than compute prices fell in the PC era, and faster than bandwidth fell in the dotcom years. The analysts who produced those figures add a caveat worth keeping: the steepest declines occurred most recently, and nobody knows whether that rate holds. The argument here does not need it to. It needs the direction, which is not in

dispute.



Sources: a16z 'LLMflation' (2024); Epoch AI inference price trends (9x-900x/yr across benchmarks, median ~50x).

Same capability, falling price. The cost of a fixed quality of analysis roughly halves every two months. GPT-3-level reading fell a thousandfold in three years; GPT-4-level is on the same slope.

Make it concrete with the document from the scene every consultation produces. A 400-page impact assessment is roughly 270,000 tokens. In 2022, no model could hold it at all; the reading was not expensive, it was *impossible*. In March 2023, models that could work through it in pieces existed, at prices and error rates that made systematic use a project. Today a long-context model reads the entire document in one pass, answers a specific question about the hydrology annex, and the cost is measured in cents, falling toward fractions of a cent.

In the economist's terms: the marginal cost of document analysis is going to zero, and the term carries the consequence. **Cheap analysis is not the same as cheap accountability.** What does not go to zero is the cost of a *proof*: a finding that stands up before a regulator or a court still requires human certification, documented method, someone accountable for the claim, and those costs fall far more slowly. What goes to zero is the analysis underneath the proof: the reading, the cross-referencing, the finding, the watching. The trajectory matters more than the level, and at roughly tenfold per year the trajectory is not in doubt.

Marginal cost near zero does not just make checking cheaper. It changes the logic of checking. When each analysis carried a price, verification was a rationing problem: you chose what to examine, and the choosing was most of the job. When each additional analysis costs approximately nothing, you stop choosing. You examine everything, continuously, as a background process; the sample becomes the census. And in that world the load-bearing question inverts: it is no longer *what does it cost to check this?* but *why has this remained unchecked?* Uncheckedness stops being the default condition of corporate claims and becomes a property that requires an explanation: indifference,

obscurity, or someone's interest in it.

The same rationing runs inside the companies doing the work, and it has a name there. Due diligence is risk-based: you segment the chain, assess, rank, and then draw a line where your capacity ends. Everything below the line is not unknown. It has been mapped, scored and set aside, with reasons, in a matrix somebody maintains. Anyone who has run one of these programmes could tell you roughly where their line sits and what is just underneath it. That line was never a statement about which risks matter. It was a statement about how much attention could be bought.

Understand what this means for companies in general, before returning to one industry in particular. The era in which a supply chain commitment could remain safely unexamined is ending, and it is ending for everyone at once. Every sustainability report ever published is becoming machine-readable against the ground it describes²⁰: the permit record, the customs data, the deformation signal, the local-language press, the court filing. The shift is from a world where companies *say* to a world where they will be expected to *show*, not because norms changed, but because the price of asking did. Firms that built their accounts for the skimming reader are holding an archive written for an examination that has now been scheduled.

One boundary belongs here rather than three parts later, because it decides how much of the above is true. What collapses is the cost of processing evidence that already exists in some retrievable form: the permit filed with an authority, the satellite pass, the customs line, the court record, the article in a local paper. It is not the cost of producing evidence that nobody has written down. A worker who is threatened, a miner working outside any register, a village whose spring changed three summers ago, a recruiter charging fees in another country: none of that generates a dataset by itself. In much of the environmental record the ground is already partly digitised, which is why those cases move first. In most labour and human rights cases it is not, and there the scarce input was never reading capacity. It was trustworthy primary information from people with reason not to give it. That inaccessibility cuts both ways: what sits outside the machine's reach is excluded from every system built on top of it, and for the same reason not yet available to be taken.


That does not weaken the argument. It locates it. Where the evidence exists, the constraint was the reading, and the reading just became cheap. Where it does not exist, the constraint is the channel through which people can safely put something on the record, and that channel is built out of trust rather than compute. Part IV is about the first case. Part VI is about the second, and the two are not alternatives.

There is a second half to this that the framing so far understates, and it is the half practitioners feel first. The same collapse that makes companies checkable also frees them. Anyone who has run supply chain due diligence knows where the time actually went: into desk research. Compiling country risk, screening suppliers, reading assessments, assembling the report: work that consumed research budgets and most of a team's attention, and that ended, in the best case, in a document. That work is becoming a background process. What a specialist provider delivered over weeks a few years ago now arrives in minutes, at a quality that holds up. And the capacity

this releases is precisely the capacity the field has always been short of: time on the ground. The sourcing team that spent most of its year assembling paper can spend it at the site, in the room with the community, inside a project that runs for years instead of a reporting cycle. The desk work collapses; the field work becomes the job. There is a sharper edge to this, too: when the thick risk report costs a keystroke, it stops counting as evidence of effort. What remains legible as diligence is exactly what cannot be generated: presence, relationships, remedy that arrived. The collapse does not only change who can check claims. It changes what counts as a serious claim at all.

Which cuts the other way too, and this is the part an optimistic reading skips. Cheap analysis produces cheap findings. Expect false positives at volume, produced by people with no capacity to check them and every incentive to publish. Expect companies generating counter-evidence at the same marginal cost, and doing it better, because they have lawyers. Expect synthetic compliance documentation, and expect small organisations to be sued for publishing a machine-assisted allegation they could not defend line by line. The economics that make accusation cheap make defence cheap and intimidation cheaper.

Follow that through and the scarce resource inverts. When evidence is abundant, the thing in short supply is credibility: whose finding is worth acting on. That favours whoever can show method, sourcing and a record of being right, which is one more reason the capability belongs with organisations that have spent decades building exactly that, rather than with whoever can afford the most inference. An accountability system built on volume alone will drown. One built on volume plus provenance will not.



The crossing. The physical curve rises, the verification curve collapses. Both are committed; only the routing is open. The series share a timeline, not a causal axis.

The two curves belong side by side, because together they are the whole argument. The same industry driving the material curve is building the collapse curve. Every model release that makes the build-out bigger makes the audit of the build-out cheaper. The AI industry is — inadvertently, structurally, irreversibly — constructing the instrument by which its own supply chains, siting decisions, water draw and sustainability claims will be examined, by anyone, at negligible cost, forever.

Put bluntly:

Every claim a company makes about its supply chain is now a liability with a delay fuse, and the fuse is shortening tenfold per year.

The equilibrium of the last thirty years (commit broadly, report confidently, rely on the improbability of the line-by-line check) was never an ethics. It was a price. The price is going away, not falling to a lower price, but ceasing to ration. What happens next is not that every company gets caught; it is that the *optimal strategy flips*. When verification was expensive, the rational posture was confident opacity. When the analysis underneath verification is marginally free, confident opacity becomes the single most dangerous posture available, because the gap between the account and the facts on the ground will be found — not by a regulator on a schedule, but by anyone, on no schedule, retroactively, across everything ever published.

So the binding constraint has moved twice in four years, and is about to move a third time: from *impossible*, to *expensive at scale*, to *marginally free*, at which point the only scarce inputs left are the questions, the truth on the ground, and the standing to act on the answers.

That last one deserves more than a clause, because it is the strongest objection to everything above. Look back at Part II. The Halmahera clearing was measured. The Morowali deaths were counted. Kachin has been documented in detail by people who took real risks to do it. The knowledge existed, and not much followed. If evidence were the binding constraint, these cases would have moved years ago.

There is a harder version of that objection, and it came from someone who has spent years doing this work at sites. If knowing something creates a duty to act on it, then whoever benefits from not knowing has no reason to fund the finding out. On that reading, the thinness of ground-level evidence is not only a resource constraint. It is partly a preference, and the field has spent years mistaking one for the other.

So the honest version of the claim is narrower than it first appears. Cheap analysis does not create consequence. What it changes is the supply of cases that are ready to be acted on, which is not nothing: most of what dies in this field dies before it reaches anyone with authority, because nobody had the weeks it took to turn an observation into a filing that a regulator, a court or a buyer's contract clause could act on. There is a difference between a harm that is known and a harm that is documented to the standard a proceeding requires, and that gap is exactly the reading gap. Closing it does not guarantee consequence. Leaving it open guarantees the absence of it.

And there is a second effect that matters more than the first. When each case cost weeks to prepare, an organisation could bring one. When it costs days, it can bring the pattern. A single site in a single valley is a local matter that can be settled locally,

usually quietly. Forty sites with the same signature, one buyer, one certification scheme, is a different object, and it is the kind of object that regulators, litigators and investors are equipped to act on. Enforcement follows aggregation more reliably than it follows severity, which is uncomfortable but observable in every sector where standards have actually moved.

One qualification, because the argument would be easy to overstate. None of this is standard practice today. Most teams in this field are at the beginning, and systematic comparison between what a company reports and what the ground shows is happening almost nowhere. This essay describes the next ten years, not the last twelve months.

Which is why the question of who ends up with the capability is still open. The tools that will be in the field by 2030 are being specified now, in decisions about which languages get measured, what sits in which tier, who sees a model early. Those decisions harden into facts about the world, and this is the short window in which they are still choices.

Who holds those is the subject of the next part.

IV. The Crossing

So the analysis underneath verification becomes marginally free. Free for whom?

The uncomfortable answer, and the hinge of everything that follows: by default, for the people who need it least.

A capability does not flow to where it is most valuable. It flows to where there is a purchase order.²¹ The first users of cheap machine reading at scale are the ones already inside the pipeline: compliance departments, ESG data vendors, law firms, consultancies. The collapse curve reaches a mining company's legal team in the same quarter the capability ships, because they have a subscription, an integration budget and a use case on file. It reaches a farmers' association above a tailings dam years later, or never. Not because anyone decided that, but because nobody decided anything. The association is not in any pipeline. It has never filed a support ticket. It generates no demand signal at all.

Run that forward and the collapse curve produces something perverse: verification gets cheap, and the *asymmetry gets worse*. Not because civil society is behind the curve. The larger organisations are already using these tools, and often well; nobody needs to explain language models to the research desk of a major human rights organisation. The difference is in what adoption looks like on each side. On the corporate side it is an enterprise contract with a budget line: the newest models, integrated into the workflow, at whatever token volume the work requires, the quarter they ship. On the other side it is a handful of subscriptions paid from project grants, stretched across an entire organisation, with no budget for integration, fine-tuning or the volumes that continuous monitoring actually takes. Both sides are adopting. One side is adopting at twice the speed with a hundred times the budget, and so the gap widens even as

everyone moves. The 54 percent of projects on Indigenous land get better-written consultation records, faster. The workers inside the fence at Morowali face an employer whose compliance documentation improves quarterly. The same technology that could close the thirty-year gap widens it.

To be fair to what already exists: the model providers have not ignored civil society. Since late 2025 the major labs have run nonprofit programmes, with discounts of seventy to seventy-five per cent on business tiers at some providers, training courses, and integrations with the platforms the sector runs on.¹⁸ These are real, and organisations are using them.

Look closely at their shape, though, because it explains why the gap described here survives them. Eligibility runs through registered charitable status, verified by a third party against an organisational email address. Discounts apply to seat licences rather than to the analysis volume that continuous monitoring consumes. The integrations connect to donor management, grant data and fundraising systems. What that adds up to is a serious offer to the professional back office of the registered nonprofit sector, and almost nothing for the actors in Part II. A farmers' association above a tailings dam has no charitable registration. A worker committee inside an industrial park has no organisational email. An Indigenous authority is not a 501(c)(3) or its equivalent, and would find the question insulting. And none of the integrations on offer connects to a permit register, a satellite archive or a body of testimony.

So the honest version of the problem is narrower than it looks and harder to dismiss. It is not that nobody is routing capability toward civil society. It is that the routing has been built around a familiar model of what a civil society organisation is: incorporated, staffed, English-speaking, with a finance system worth integrating. The organisations holding the truth on the ground in these supply chains are mostly none of those things, and the gap that remains is one of specification rather than generosity.

Unless someone routes it. And this is where civil society stops being a stakeholder paragraph and becomes the strategic centre of the whole picture.

Civil society is not a beneficiary of the collapse curve. It is the mechanism through which the collapse curve becomes real.

Consider what actually has to happen for cheap analysis to change anything. Someone has to know which permit matters and which paragraph of it. Someone has to be trusted by the people downstream of the dam, because the truth on the ground lives with the people who see the creek every day, and they do not hand their observations to a dashboard. Someone has to be able to reach the workers inside the industrial park, which the accident count from Part II shows no announced audit ever will. Someone has to carry a finding to the authority that can act on it, in the form that authority accepts, and stay on it for the four years that takes. Someone has to be believed when they publish.

None of that is a model capability. All of it already exists, distributed across NGOs, community and Indigenous organisations, legal aid networks, unions, investigative desks, church networks. Thirty years of documented harm, filed testimony, local knowledge and hard-won trust: that is the question bank and the distribution network

of accountability, and it is held almost entirely by organisations running on three people and project grants. What they lack is exactly and only the thing whose marginal cost just went to zero: the capacity to read, cross-reference, and process at scale.

Put the two halves together and you get the real picture of the crossing. The model provides the reading. Civil society provides everything that makes reading matter: the questions, the truth on the ground, the trust, the follow-through. Neither side can produce verification alone.

The collapse curve plus civil society is an accountability infrastructure; the collapse curve minus civil society is a compliance product.

It is worth seeing what that looks like in practice, and the rivers from Part II are the clearest available example. Everything needed to make that case already exists, held in six different places. A national pollution authority has its monitoring series. A university toxicologist has samples from seven points, many of them collected and handed over by the people who live there. A research institute abroad has satellite imagery showing when each mine opened and where. Customs data records what crossed into China and when. Residents have years of observation about which season the fish went. A newsroom has costed the damage to twenty communities.

No two of those datasets have been joined into the thing that would actually move: a chain of attribution linking specific operations to specific harm, in a form a proceeding could act on. The best public estimate attributes sixty to seventy per cent of the river's contamination to rare earth mining, which is a start and not a case. Nobody is withholding anything. The pieces sit in six institutions with different formats, different languages and no shared mandate, and the work of joining them is reading. That work was, until recently, more expensive than any of the six could justify.

Now a human rights case, and this one matters more because it sits where audits structurally cannot reach. Workers at a processing site pay a recruitment fee before they ever arrive, in their home region, to an agent who has no contractual relationship with anyone in the supply chain. An announced audit will not find it. The workers are on site, in working hours, within earshot of the employer, and the harm happened months earlier in another country. But the workers know. If there is a channel they trust, they can say so in their own language, over a phone, outside the gate. Individually each account is an anecdote. Forty of them, structured and compared, show the same agent, the same fee, the same three labour suppliers. That pattern is evidence, and it is evidence no audit cycle would ever have produced, because the cycle was never looking there.

Note who benefits. The workers get a route to remedy. The buyer, who has a contractual clause about recruitment fees and no way to check it, finally gets something to act on. Even the operator learns about a subcontractor it may genuinely not have known was in its chain. This is the non-zero-sum part, and Part V returns to it. What makes it work is not the model. It is that the people with the knowledge had somewhere to put it and someone who could act on it, with the reading step in between no longer costing more than anyone could pay.

And note the two places in Part II where this is not merely the better option but the

only one. In the Kachin hills, there is no operator to audit and no site to visit; remote analysis joined to the organisations that hold the knowledge on the ground is the entire universe of possible due diligence. In the midstream parks, the truth on the ground is held by the workforce, and only channels that workers control, protected and outside the employment relationship, can carry it. The two least-governed links of the chain are precisely the two where the civil-society route is not the alternative to the corporate toolkit. It is the toolkit.

Before drawing out what this demands of the companies, one more piece of the picture, because the claim so far could be misread as saying nobody ever built accountability right, and that is not true.

V. What already exists, and what it has been waiting for

Before the responsibility question, an honest correction to the picture so far. Part II described what the accountability machinery lacks. It did not say that nobody has built the right architecture. Somebody has.

The Initiative for Responsible Mining Assurance, IRMA for short, is worth a paragraph of detail, because it inverts the design flaw that runs through the assurance industry. Its governance does not belong to the companies being assessed. Six sectors govern it with equal authority: mining companies, downstream purchasers, finance, affected communities, organised labour and NGOs. Each holds two seats on the board, and each has veto power, since two no votes from one sector block a decision. Mining is one voice of six, not the house. Its standard was written in that room, over years, with the people who live next to the mines at the table rather than in an annex. Its audits are independent, its site reports are published in unusual detail, and, rarest of all, affected people can trigger and participate in the assessment of the mine next to them. Measured against the ten-year study's finding that fewer than one in five initiatives include rights holders in governance at all, IRMA is the counterexample: proof that the architecture is buildable, because it has been built.¹⁴

And now look at what constrains it, not as failure, but as the outline of what needs building out. An IRMA audit is still a *visit*: a snapshot every few years, at a facility that knew the date. Between visits, the standard lives on paper while the water table, the embankment and the night shift live in the world. Its reach is still a fraction of the sector; the mines most in need of assessment are the ones least likely to volunteer for it, and nothing in a voluntary standard reaches Kachin. Its measurement, like everyone's, can say far more about process than about what changed for the people downstream. And the communities inside its governance carry the same capacity constraint as everywhere else in this essay: a seat at the table, and three staff to read four hundred pages before the meeting.

The same pattern shows up in a place the previous part treated harshly, which is corporate practice itself. Alongside the reporting there is field engagement: buyers who put people at a mine or a processing site for years rather than weeks, who fund

remediation without a certificate at the end of it, who stay through the part where progress stops being reportable. Where it happens, it produces the one thing no assurance system reliably produces, which is somebody who actually knows the place. And it is the most expensive form of diligence there is. A single site can absorb a team's attention for years, which is why this work exists as individual projects rather than as programmes, why it tends to concentrate on a handful of high-visibility supply chains, and why the same company's engagement can be deep at one site and paper-thin across the rest of its footprint. Nobody chose that unevenness. It is what a budget for presence buys when presence is priced by the day.

The same architecture exists at the informal end of the sector, and the same kind of constraint binds it. The ASM standards, Fairmined and the CRAFT code among them, were written with miners rather than about them. Around them sits some of the most serious field practice in sourcing: the Alliance for Responsible Mining, IMPACT with its traceable gold work in the DRC, Pact, Solidaridad and others, embedded at sites for years, building formalisation step by step. Helping miners organise into cooperatives, obtain licences, meet the entry requirements formal buyers demand, and sell into legal channels at fair terms. Where these programmes run, they work. What they have never been able to do is spread. A single site can take years of presence, trust-building and mediation before the first compliant shipment leaves, and the informality that makes ASM hard to source from is what makes every one of those years expensive. Twenty years in, the field is full of good pilots and short of programmes, and the projects that succeed cover a rounding error of forty-five million people.

Part of what makes each site slow is a wall of paper. Formalisation runs on documents: licence applications, land and title records, cooperative registration, the paperwork buyers require before they touch a shipment. The binding constraint in ASM is income and presence, and cheap analysis feeds no one. But the document wall is one reason each site takes years, and it is exactly the kind of wall that is getting cheap to climb. Cheaper engagement per site does not solve ASM. It changes how many sites the existing organisations can be present at, which is the difference between a pilot and a programme.

What do all these constraints have in common, from IRMA's audit cycle through the cost of a team on site to the years behind one compliant ASM shipment? Continuous evidence instead of snapshots. Ground truth from the people who are there every day, carried through protected channels. Reading capacity on the affected side of the table. The cost of standing up one more site. Every one of them is, at bottom, a function of the cost of analysis and engagement. And that cost is the thing Part III watched go to zero.

The best-built architectures in the sector have been waiting, without knowing it, for exactly the price collapse the AI industry just delivered.

With that in view, the responsibility question can finally be asked properly.

VI. The Responsibility

What, then, do AI companies owe here? Not as philanthropy, a framing that should die because programmes conceived as charity are scoped as charity, budgeted as charity and cancelled as charity, but as responsibility: something that follows from what these companies are and do, and that can be argued about, demanded, and checked.

One distinction first, since the rest depends on it. The framework establishes that a company connected to harm through its business relationships carries a responsibility, and that it should use and build leverage rather than plead the absence of control. That much is the standard, and it is not mine. What follows from it here is not. The five components below are my reading of what that responsibility implies for a company whose product happens to be the analytical capability in question, and they should be argued with on those terms rather than treated as obligations the framework already spells out. It spells out the duty. It does not spell out the evaluation suite.

The claim of this essay is that the responsibility has five parts. They apply to any company building frontier models; they scale with the size of the build-out; and none of them requires believing anything except Parts I through V.

First: publish your own numbers

The starting point is a no-brainer, and it should read like one. AI companies are connected, through their business relationships, to everything described in Part II.⁷ That makes them ordinary companies in these chains, and the first duty of an ordinary company in these chains is disclosure: publish an honest account of your own position. Not a programme, not a pledge. Numbers.

Which numbers, if you own the infrastructure. Absolute material quantities, because intensities per query conceal growth. Water by site with the local stress context, because an aggregate hides exactly the sites that matter. Hardware refresh cycles and end-of-life ownership, because lifetime is the largest material lever a company controls entirely on its own. Power sourcing matched by the hour, because annual averages say little about what was actually running when. None of this needs a new discipline; responsible sourcing has done it for twenty years. What is new is the volume, and that someone will check.

If you buy compute rather than build it, the honest version is smaller and still worth having. Publish what you know about the infrastructure your models run on, what you asked for contractually, and where your visibility ends. A clean material position that a company cannot actually see is the kind of claim Part III is about. Saying plainly how far your sight reaches, and what you did with the leverage a large compute contract carries, is a claim that holds.

One part of the account takes more nerve than the rest: the classification the responsibility framework itself asks for. Where does the company cause harm, where does it contribute, where is it linked. Almost nobody publishes the middle column, because contributing carries a duty to help fix what you contributed to. That is exactly why publishing it is the one honesty signal that cannot be faked, and why

deteriorations belong in the account as much as improvements.

Part III explains the urgency. This account will be read. Not skimmed by a rating agency, but checked line by line against the ground it describes, by anyone, at negligible cost, retroactively. The only sensible move is to write it for that reader from day one. A company that ships verification capability to the world while its own numbers stay unpublished has answered the seriousness question before anyone asked.

Second: treat capability routing as the decision it is

This is the responsibility that belongs to AI companies alone, and it is the heart of the matter.

Underneath every application anyone builds on a model sit choices the model provider has already made: which languages work and which do not; what sits in the base model and what behind the enterprise tier; what is measured in the public evaluation suite and therefore improves; how long a model version survives before deprecation kills everything built on it. These choices are currently made by default, for unrelated and usually good reasons, by small teams optimising for other things, and Part IV showed what the default does: it routes the collapse curve entirely to the side of every dispute that has a purchase order.

The responsibility is not to build any particular thing. It is to *make these decisions looking*, knowing that a language left unmeasured is a population left outside the record, that a tiering choice is a distribution choice, that a deprecation schedule decides whether anything maintained by one volunteer survives its first year. A company can weigh those consequences against its other constraints and decide differently in good faith. What it cannot do, once the routing effect is named, is claim the outcome was nobody's decision.

Third: equip the side of the table that cannot pay

Part V showed the best architecture in the sector constrained by reading capacity on the affected side. Part IV showed that civil society, communities, Indigenous organisations and unions are not beneficiaries of the collapse curve but the mechanism through which it becomes accountability at all: the holders of the questions, the truth on the ground, the trust and the follow-through. And Part II showed two links of the chain, Kachin and the midstream, where their route is not the alternative to the corporate toolkit but the only toolkit there is.

The responsibility that follows: ensure the capability reaches that side of the table on terms it can survive on. That means the affected side helps decide what gets built, in arrangements it co-designs. It means paid participation and working languages that are theirs. It means funding the actors with the better edge instead of duplicating them. And it means measured, published performance in the languages that matter, speech included, because the people above the leaching ponds speak and do not necessarily write.

One correction to the picture this essay has drawn so far: the side of the table

that cannot pay is not only human. Rivers, wetlands, forests and the species in them are parties to every dispute in Part II, and they file no objections at all. Their case is carried by proxies, by conservation organisations, environmental defenders, community monitors and scientists, and those proxies face the same capacity wall as everyone else on that side. The difference is that here the raw evidence is unusually abundant and much of it is free: public satellite programmes image every mine on earth every few days, water records exist, acoustic and camera-trap datasets exist. What has been scarce is the work of turning observation into a case. Connecting the pixel change to the permit condition, the water series to the abstraction right, the species record to the licence condition that promised otherwise. That connection work is document-and-data analysis, which places it squarely on the collapse curve. Equipping environmental advocates is therefore the same responsibility as equipping communities, applied to the rights holder that cannot speak at all.

Where a company's own product decisions cannot deliver this, its scale can. The sums involved round to zero against a single campus, which removes the only respectable excuse.

Fourth: worker voice, the clearest case and the sharpest line

Of everything in this essay, the midstream gap from Part II is where the mechanics deserve to be spelled out, because it is the place where what AI changes is most concrete, and where getting the architecture wrong does the most damage.

Recall the problem's shape. The truth about a smelter is held by the eighty thousand people inside it. The audit interviews a sample of them on site, in working hours, within earshot of the employment relationship, and learns what that arrangement permits it to learn. The model that actually works runs the other way: worker-driven monitoring, where workers themselves report through a channel they trust, individual reports are protected, and consequences flow through buyers with leverage. It exists: the Fair Food Program built it for US agriculture over a decade, around worker-to-worker education and a complaint line backed by market enforcement, and it is widely regarded as the most effective labour-rights mechanism of its generation, cited as such by the UN Working Group on Business and Human Rights.¹⁶ What has kept it from the midstream of the AI bill is not the design. It is the operating cost of the front end: intake in the languages a workforce actually speaks, triage at the volume a real channel produces, feedback to every person who called. All of it specialist attention, priced accordingly.

That front end is exactly what just collapsed in price, though not in the way the phrase suggests. The intake itself stays human. The evidence on automated questioning of people describing harm is unambiguous, and the established standards for interviewing survivors exist for reasons that no model satisfies: trained interviewers, referral pathways agreed before anyone is asked anything, the capacity to notice when a conversation should stop. A system that interviews at scale is not the goal and would not be an improvement.

What the model does is everything after. Testimony arrives *spoken*, in Bahasa or

Burmese or Hausa, over a phone or in person, outside the gate and outside working hours, taken by someone from an organisation the caller already trusts. No form, no literacy requirement, no employer-issued app. The model then transcribes, translates and structures what was said, and a human at the receiving organisation certifies before anything moves. Individual voices never surface. What surfaces are *patterns*: forty separate accounts of the same furnace, the same unpaid hours, the same retaliation after the same safety complaint, aggregated past the point where any single speaker can be identified, which is the transformation that turns anecdote into evidence while protecting the person who provided it. The pattern goes to a recipient with the power and the standing to act: the union, the labour inspectorate, the buyer whose contract carries a labour clause, the worker-driven monitoring body. And the answer travels back down the same channel, in the caller's language, because a channel that only takes is a channel people rightly stop using.

Now the responsibility question, and it has an unusually clean answer in three steps. *Only AI companies decide whether the front end exists*: whether speech recognition works in the languages of the industrial parks is set in training and evaluation choices nobody else can make, and today those languages are mostly unmeasured, which is a decision by omission. *The AI company cannot run the channel itself*: a commercial party, plausibly a supplier or customer of the very operators concerned, cannot credibly hold worker testimony — trust is the one component that must come from unions, NGOs and worker organisations, who have spent decades earning it. So the responsibility takes the only shape that works: build and measure the capability, and put it in the custody of the organisations workers already trust, with published per-language performance, support horizons their one technical volunteer can survive, and funding where product access is not enough. This is not outsourcing the duty. It is the only architecture in which the duty can be discharged at all.

One threshold is worth stating, because it may move. If open speech models reach dependable accuracy in the languages that matter here, and if evidence emerges that automated intake can be done without harming the people it is meant to serve, the balance shifts. Neither has happened. Until both do, the rule is people first.

And here the line from Part IV must be drawn at its sharpest, because the same capability has a mirror image. Speech analysis sold to the employer's side of the fence is sentiment monitoring, complaint triage before the complaint leaves the building, organising detected early: the surveillance layer of every labour dispute, marketed as engagement analytics, and it is where the purchase orders are. A company that routes voice capability only to that side has not stayed neutral; it has armed the fence. So the uncrossables stand: testimony belongs to the people who risked giving it, sits with neither the AI company nor the operator, and feeds nothing else. No deployment without a capable recipient; a channel that collects and cannot deliver exposes people for nothing. And the safety of the people speaking is a design constraint with veto power, not a review item. A company that gets everything else in this essay right and this wrong has built the instrument of the other side and called it accountability.

Fifth: be examined, do not self-report

The last responsibility closes the loop with Part II's finding about the account and the ground. A company acting on the first four cannot be the narrator of its own success; that is the equilibrium this essay watched die. The external arrangement writes the assessment and publishes it unedited; the company replies in an annex. What was asked for and what was declined is public, with reasons. Failures and discontinuations appear in the record, because a portfolio in which nothing ever fails is being administered, not measured. The one asset the collapse curve makes valuable, claims that have survived hostile reading, can only be accumulated this way, and it cannot be bought.

Five responsibilities. None requires a strategy cycle; a company that accepted them could show verifiable movement on every one within a year, at a cost that disappears against the capital expenditure of a single campus. And they compress, for anyone evaluating what a company announces in this space, into three questions.

Who decides what gets built? Who holds the grievance and worker data? Who writes the assessment?

Everything else is detail.

The objection from Part IV deserves its final answer here: yes, the same capability makes the companies' own defence cheaper too; both sides get armed either way. That is exactly why the responsibility is framed as routing rather than building. The default arms one side completely and the other not at all, and the default requires no decision. These five require one.

VII. The Decade

Pull the threads together.

The material curve is committed. The gigawatts are announced, the transformers are ordered, the copper will be mined, the magnets will come from wherever the magnets come from; nothing in this essay slows a single campus. The collapse curve is equally committed; no one can decide that reading stays expensive. The shift from claiming to proving will happen to every company with a supply chain, whether or not anyone plans for it. The only open variable, the only place where a decision exists at all, is the routing: whether the analytical capability lands solely where the purchase orders are, or also where the truth on the ground is: with the communities on the 54 percent, the workers inside the parks, the organisations that documented Kachin when no one else could get in.

That variable gets set in the next few years, mostly implicitly, in evaluation suites and tiering decisions and depreciation schedules, by a few thousand people at four or five companies who largely do not know they are setting it.

Neither position has to wait for the other. The infrastructure owners have the

sourcing functions and the supplier relationships already, and have not connected them to the question of who gets to use the capability. The model providers hold that question outright. Where both sit in one company, one decision covers both.

Risk-based due diligence does not disappear in any of this. Severity still governs, and you still start with the worst. What changes is what the line means. Until now, “below our prioritisation” was a complete answer, and a legitimate one, because capacity was genuinely scarce. As analysis stops being the scarce part, it turns into a question instead: why is your line still there? Not every constraint lifts at the same rate. Presence does not get cheaper, and some of what sits below the line needs a person on site rather than a better reading of a document. But knowing which cases those are is itself the thing that was missing.

The more consequential shift is in timing. Diligence has been a cycle because a pass through the chain was expensive: assess, report, wait a year. Cheap analysis plus a channel the affected side actually controls turns that cycle into something closer to observation, and observation shows a different class of thing. Not incidents but trajectories. A well drawing down over two seasons. The same recruitment fee surfacing at three suppliers. An embankment deforming by millimetres. These are the harms that currently become visible as catastrophes, because between audits nobody was looking. Seeing them earlier is not a smaller version of the same accountability. It is the difference between remedy and compensation.

None of that follows automatically from cheaper models. It follows only if the people living with the operation are inside the channel, because they are the ones who notice the well before the satellite does. Get that wrong and the same capability produces faster, better-documented reporting about nothing in particular.

One absence in this essay is deliberate and should be named rather than left to be noticed. It says almost nothing about the law: the European due diligence and reporting regimes, their national predecessors, or the political fight over how much of them survives. That is not because the law is a side issue. It is because the argument here is about a constraint that operates underneath whichever rules apply, and it points in a different direction depending on how that fight ends. If mandatory duties are scaled back, cheap verification becomes the main check that remains, and it will be exercised by the public rather than by supervisory authorities. If they hold, it becomes the thing that makes them enforceable, since a duty to identify risks means little while identifying them is priced beyond most of the parties involved. Either way the collapse is happening, and either way the routing question is the one that is still open. What the law does determines how much the answer is worth, not whether it has to be given.

Two futures, then, both entirely plausible.

In one, the default runs. Verification capability industrialises on the compliance side; the account of every supply chain gets better-written faster than it gets better-checked; civil society adopts on grant budgets what the compliance side integrates on enterprise contracts. Sometime in the window, someone outside runs the check that has become trivial, against a company still writing for the old equilibrium, and publishes the gap. The industry gets its demonstration case, the story about broadly beneficial technology

takes the hit it will by then deserve, and the regulation that follows is written by people holding the receipt.

In the other, at least one company reads the curves and moves first. It publishes its own bill before anyone demands it. It routes reading capability to the people the framework says should have been consulted all along. It lets outsiders set the agenda and write the verdict, and it publishes the list of what it declined to build. It is examined constantly — that is the point — and its claims are the only ones in the industry that have survived hostile reading, which by then is the only kind there is. It has not solved water in the Atacama, or the leaching ponds in Kachin, or the furnaces in Morowali; nothing in a model does. It has done something narrower: ended, for its own chains, the era in which the account and the ground could drift apart unexamined, and set the standard every competitor is now measured against, at a cost that rounded to zero.

The first company to do this will get it credited as conviction; the second will get it discounted as response.

For thirty years, the distance between what companies said about their supply chains and what was true of them was protected by the price of reading. The price of reading is going to zero, on a curve steeper than any in the history of this technology, built by this industry, billed to it either way.

The only question left is who notices first.

This essay is personal work. It reflects my own analysis and no one else's, and it is not connected to my employer or to any organisation I work with. I wrote it because the two curves it describes are moving whether or not anyone discusses them, and because the decisions that follow are still open. If it prompts an argument about where the capability should go, it will have done its job.

Written with AI-assisted research and drafting; the argument and its errors are the author's.

Sources

1. IEA, *Energy and AI* (2025): 415 TWh (2024) → ~945 TWh (2030) base case; datacentre demand growing four times faster than all other sectors combined.
2. IEA (installed capacity ~100 GW); McKinsey & Company (2025): AI workloads ~156 GW by 2030.
3. Meta Hyperion campus, Louisiana: planned build-out to ≥ 5 GW. The comparison to New Orleans refers to that city's annual electricity consumption, not to continuous load (Paul Arbaje, Union of Concerned Scientists, via CBS News, 2025). Other published comparisons range from three to seven times, depending on whether annual consumption or peak load is used.
4. IEA, *Energy and AI* (2025): up to 20% of planned datacentre projects at risk of delay on transmission constraints.

5. Bloomberg Intelligence: 27–33 tonnes of copper per MW of applied power.
6. CRU / Wood Mackenzie estimates (2025): ~700 kt copper in datacentre facilities through 2030; up to 5 Mt in associated T&D build-out, ~1.1 Mt directly attributable. Commercial estimates; ranges vary by methodology.
7. UN Guiding Principles on Business and Human Rights (2011): Principle 13(b); commentary to Principle 19 (leverage: use, build, disclose; consequences of remaining); Principles 18, 21, 24, 31.
8. Siddik, M.A.B., Shehabi, A. & Marston, L. (2021). The environmental footprint of data centers in the United States. *Environmental Research Letters* 16(6), 064017.
9. Global Tailings Portal, GRID-Arendal for the Investor Mining and Tailings Safety Initiative, public since January 2020: ~1,800 facilities from voluntary disclosure by 726 companies approached. Brumadinho failure, January 2019: 272 deaths. GISTM: ICMM/UNEP/PRI, 2020.
10. Owen, J.R., Kemp, D., Lechner, A.M., Harris, J., Zhang, R. & Lèbre, É. (2023). Energy transition minerals and their intersection with land-connected peoples. *Nature Sustainability* 6, 203–211. 5,097 projects; 54% on or near Indigenous peoples' lands; >50% of the resource base on or near Indigenous or peasant lands; lithium 85%, manganese ~75% overlap with Indigenous lands.
11. World Bank / Pact, *State of the Artisanal and Small-Scale Mining Sector* (Delve); World Bank (2025): ~45 million directly employed across ~80 countries; ~315 million including indirect employment. On field practice and formalisation: Alliance for Responsible Mining (Fairmined standard; CRAFT code, with RESOLVE), IMPACT (traceable artisanal gold, DRC), Pact, Solidaridad.
12. Global Witness, *Fuelling the future, poisoning the present: Myanmar's rare earth boom* (2024); ISP-Myanmar (2025); Mongabay / Myanmar Witness satellite analyses (2025–26): ~370–400 mining sites and ~2,700–3,000 leaching ponds in Kachin State; >290 kt exported to China 2017–2024 (~two-thirds of China's rare earth imports); Myanmar the world's largest heavy rare earth source since 2021; no environmental impact assessment ever conducted for any rare earth operation in the country (*The Lancet Planetary Health*, 2026). On the downstream effects in northern Thailand: sampling by Wan Wiriya, Chiang Mai University (seven points on the Kok river); Thai Pollution Control Department monitoring rounds to August 2025 against a 0.01 mg/l arsenic standard; satellite analysis of mine sites by the Stimson Center (2025); economic damage estimate for twenty communities in Chiang Rai by the independent newsroom Lanner; reporting by Tagesschau (Christiane Justus, 16 March 2026) and Mongabay (2025). The attribution of 60–70% of river contamination to rare earth mining is an expert estimate by Tanapon Phenrat (Thailand Science Research and Innovation), not a forensic finding.
13. Business & Human Rights Resource Centre / Trend Asia / Financial Times reporting (2023–25) on Indonesian nickel processing: 114 incidents, 101 deaths, 240 injuries 2015–H1 2024 (Trend Asia count from media reports; no official register exists); ITSS furnace explosion at the Indonesia Morowali Industrial

Park, 24 Dec 2023, 21 deaths; documented retaliation against workers raising safety concerns; >24 workers across nine companies interviewed, three unions representing ~12,000 members; SBIPE union founded Feb 2024. On nickel-driven deforestation in Halmahera: Climate Rights International with the University of California, Berkeley AI Climate Initiative, *Nickel Unearthed* (2024) and *Ongoing Harms, Limited Accountability* (2025), geospatial analysis identifying at least 5,331 ha cleared within nickel concessions; Forest Watch Indonesia; Mongabay reporting;.

14. MSI Integrity (2020), *Not Fit-for-Purpose*: of 40 standard-setting multi-stakeholder initiatives, fewer than 1 in 5 include rights holders in governance; 5 of the 20 oldest had measured impact on rights holders in the preceding five years. Contrast: Initiative for Responsible Mining Assurance (IRMA), governed by a board of two representatives from each of six sectors (mining companies, downstream purchasers, finance, affected communities, organised labour, NGOs) with equal authority and sector veto, independent third-party audits, and public site reports. On the scorecard generation: Lead the Charge, Auto Supply Chain Leaderboard (4th edition, 2026), coalition-run annual ranking of 18 automakers based on publicly available company reporting, including an assessment of third-party auditing and certification schemes; similarly KnowTheChain and the Corporate Human Rights Benchmark.
15. Harrison, J. & Wielga, M. (2021), *Business and Human Rights Journal*: RSPO complaints mechanism performs well against procedural effectiveness criteria, poorly on case outcomes for rights holders.
16. Worker-driven Social Responsibility: the Fair Food Program (Coalition of Immokalee Workers, since 2011) combines worker-to-worker education on company time, a 24/7 worker-triggered complaint line, independent monitoring, and legally binding agreements with buyers that carry market consequences. The UN Working Group on Business and Human Rights cited it in 2021 as demonstrating that the usual shortfalls of site-level grievance mechanisms can be avoided where workers lead in shaping and monitoring them. The model has been replicated in the Bangladesh Accord, the Lesotho Agreements and Milk with Dignity.
17. a16z, “LLMflation” (2024); Epoch AI, LLM inference price trends: 9x–900x per year across benchmarks, median ~50x; GPT-4-class from ~\$30/M tokens (03/2023) to <\$0.50; ~1,000x over three years one tier down.
18. Nonprofit access programmes launched from December 2025: Anthropic (Claude for Nonprofits, 70–75% discount on Team and Enterprise plans, integrations with Benevity, Blackbaud and Candid, eligibility via registered charitable status verified by a third party), Google (comparable discounts), OpenAI (20–25%). Programme terms change frequently; figures as reported in December 2025 and mid-2026.
19. On the extractive supply chains of AI as a field of study: Crawford, K. (2021).

Atlas of AI. Yale University Press; Crawford, K. & Joler, V. (2018), “Anatomy of an AI System”. On quantified mineral demand: FP Analytics, “Artificial Intelligence and the Critical Minerals Crunch” (2025).

20. On machine-assisted verification of sustainability claims: Bingler, Kraus, Leippold & Webersinke (2022), “Cheap Talk and Cherry-Picking: What ClimateBERT has to say on Corporate Climate Risk Disclosures”, *Finance Research Letters*; and the AIMS.au / AIMSCheck work on automated review of mandated modern slavery statements (ACL, 2025). Neither supports the specific rate of price decline used here, which is drawn from source 17.
21. On the distribution of AI capability as a governance question: Sastry, Heim, Belfield et al. (2024), “Computing Power and the Governance of AI”; Institute for AI Policy and Strategy, “Asymmetry by Design: Boosting Cyber Defenders with Differential Access to AI” (2025). On the limits of disclosure as an instrument: Fung, Graham & Weil (2007), *Full Disclosure: The Perils and Promise of Transparency*, Cambridge University Press.

Commercial analyst estimates (2, 5, 6) are order-of-magnitude anchors rather than precision claims, and forecasts in this area vary considerably by methodology. The incident counts in 13 are civil-society compilations assembled in the absence of any official register.