No Exchange Rate

Carbon in the Economics of AI Compute

James E. Murphy

Working draft for discussion · v1.0 · 1 August 2026

Abstract. The emerging discipline of AI unit economics has identified a foundational measurement question: should computational efficiency be governed by intelligence per watt or by total cost of ownership per intelligence? This paper argues that the question identifies not a choice between competing metrics but the need for a broader measurement architecture. Financial cost, energy demand, and greenhouse gas emissions are distinct but interacting dimensions of AI production. Energy is monetized through electricity and infrastructure prices, while emissions depend on the timing, location, and supply conditions of consumption, and converting them into dollars requires a carbon value whose meaning varies across compliance systems, voluntary claims, policy appraisal, internal planning, and abatement decisions. No single value is suitable as a universal conversion factor, and a monetization is well formed only when its value, its emissions basis, and its decision purpose are mutually compatible. The appropriate architecture is a vector of financial, energy, and emissions measures, each normalized to a defined outcome, with contractual instruments registered beside the physical accounts rather than netted into them, reported figures derived from both under named rules, and with the decision rule made explicit through constraints, Pareto analysis, or shadow prices specific to a scenario. Proposed changes to electricity emissions accounting illustrate why the separation matters operationally: identical computational work can carry different financial, physical, and contractual attributes depending on when and where it runs. The result offers an extension to the emerging Tokenomics framework: the field's intelligence per watt and TCO per intelligence are both necessary, neither is sufficient, and the discipline's next task is to standardize how they operate alongside carbon.

1. The question the field asked

The economics of AI compute is becoming a discipline for the ordinary reason disciplines form: the industry now needs common units, common boundaries, and common decision rules for a spend it can no longer treat as an engineering detail. The founding materials of that discipline define its subject as the conversion of energy and capital into tokens, tokens into intelligence, and intelligence into value, with power the binding constraint at scale and cost spread across nine layers of which only one, the token bill, is metered (Linux Foundation 2026; Tokenomics Foundation 2026). And among the open questions the field's founders committed to resolving, one stands out as the central measurement problem: whether the efficiency of that conversion should be governed by intelligence per watt or by total cost of ownership per intelligence, given that sustainability and cost efficiency often pull in different directions (FinOps Foundation 2026, opening keynote). The discipline's first published framework arrived in July 2026: Big-T notation, which classifies workloads by how token consumption grows as usage scales, with classes compared at a stated quality floor (Neff 2026). The question asks, in effect, whether AI efficiency is fundamentally physical or financial. The strongest answer is that the production system is both, and that its environmental consequence introduces a third dimension neither metric captures alone. This is not an objection to the question; it is an attempt to complete it.

Completing it begins with a clarification the field's own vocabulary invites. Intelligence per watt is an energy metric, and energy already converts to dollars: through tariffs, through power purchase agreements, through the energy component buried in cloud pricing. If the two founding metrics differed only in whether energy appears physically or financially, the choice between them would be a reporting preference. What makes the question hard is a third quantity that enters through energy but is not energy: the greenhouse gas consequence of when and where the electricity is drawn. Total cost of ownership conceals the physical footprint entirely; energy efficiency reveals the physical footprint but conceals cost and, by itself, says nothing about emissions, because the same kilowatt-hour carries different carbon depending on the hour and the grid. Neither metric is wrong. Neither is sufficient. The object being measured is vector-valued.

That is the paper's thesis in one image: the two metrics are instruments on one control panel, and the environmental dimension is a third instrument that belongs beside them. Read this way, the founding question converts from an adjudication into a standardization agenda, which is work a discipline can execute. It also marks a recognizable moment: the point at which unit economics becomes multidimensional is what it looks like when a discipline outgrows a scalar.

Stated as a contribution: AI compute efficiency should be represented as a vector of financial cost, energy and capacity demand, and emissions, each normalized to a defined outcome, accompanied by a disclosed decision rule for resolving the trade-offs among them. Concretely, the paper develops a common outcome denominator with fixed quality and service constraints, a measurement schema separating cost, energy and capacity, and physical emissions, an instrument register keeping contractual claims structurally apart from measurements, derived reporting results treated as outputs of named rules rather than primitives, required boundary disclosures for each component, a taxonomy of decision rules, a compatibility condition requiring emissions basis, carbon value, and decision purpose to align, and a minimum reporting profile from which an interoperability standard can be developed, making the field's existing measures composable. Figure 1 draws the architecture the contribution names.

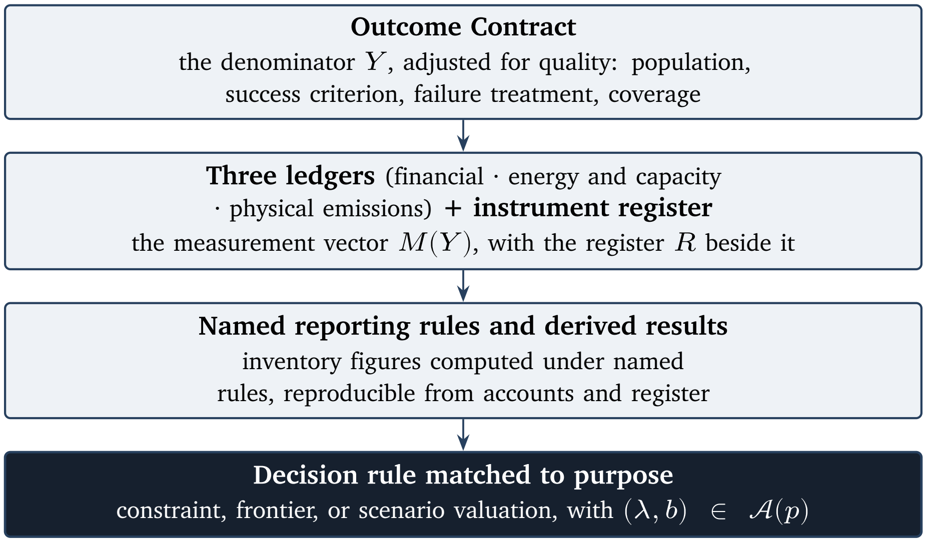

Figure 1. The architecture in one image. From denominator to decision: measurement fills the vector, named rules derive the reported figures, and governance, not arithmetic, produces any scalar, with basis, value, and purpose mutually compatible.

The stakes are not academic. Industry projections put data center investment driven by AI in the trillions of dollars this decade (McKinsey & Company 2025), with the major operators' capital spending tied to AI estimated near half a trillion dollars in 2026 alone (Goldman Sachs Research 2026a, 2026b). Data centers are estimated at roughly four and a half percent of US electricity today (Shehabi et al. 2024; EPRI 2024), with official and industry forecasts placing them between seven and seventeen percent by the end of the decade (Shehabi et al. 2024; E&E News 2026). And inference, the production phase, recurs with every served request by construction, which means the measurement question attaches to a recurring operating flow, not a build done once. A discipline standardizing its measures now will fix, for that flow, what gets counted and what gets collapsed. The rest of this paper is about what should stay uncollapsed, and why.

The argument proceeds in seven steps. Section 2 lays out the dimensional structure of AI compute: what is measured, in what units, and why the measures do not reduce to one activity base with different rates. Section 3 establishes the central negative result: the available carbon values are different decision instruments, each meaningful within its purpose, none of them a universal conversion factor. Section 4 grounds that result in a live standards process, the revision of electricity emissions accounting, where different legitimate questions visibly produce different numbers about the same consumption. Section 5 turns constructive: three decision architectures, all assembled from instruments already in use, and the metric stack they add up to. Section 6 classifies how an operator holds position while the measurement rules themselves move. Section 7 assembles the deliverable: the interoperability rules, the Outcome Contract, and the Minimum AI Compute Impact Reporting Profile with a completed record. Section 8 states what the answer changes, and ends where the argument must.

2. The dimensional structure of AI compute

Completing the field's question requires the object's dimensions on the table. Five quantities do the work, and they must be defined and then held apart.

The first is useful output, written here as Y: a quality-adjusted outcome specific to the task, a resolved support ticket at a stated containment threshold, a document processed to a stated accuracy, a code change accepted under stated review. The second is computational work, the tokens, operations, or requests the system executes to produce Y. The third is energy, in kilowatt-hours or joules, which must be distinguished from power, in watts: energy is a flow consumed over time, power is the instantaneous draw and the capacity that must be provisioned for it. The field's own headline metric respects the distinction more than its shorthand suggests, since tokens per second per watt, integrated over a workload, is tokens per joule. The fourth is emissions, in CO2-equivalent. The fifth is financial cost, in dollars.

Emissions relate to energy as follows, and Section 4 lives inside these definitions:

$$G_A = \sum_{r,h} K_{r,h} \, f^{A}_{r,h} \qquad\qquad G_{\text{total}} = G_A + G_E$$

The attributional operational account G_A sums energy consumed over regions and hours against the applicable attributional factor $f^{A}_{r,h}$, which depends on where and when the energy is drawn; total inventory adds the embodied account G_E, which arrives amortized inside hardware and facility. A consequential estimate G_Q, the change in system emissions caused by the workload's existence, is a different causal object, defined against a counterfactual rather than by an allocation, and is never summed with the others; the distinction is marked here once and spent in Section 4.

A bracket bounds every claim that follows. Tokens are not intelligence, and this paper assumes no universal unit of intelligence exists. The analysis applies within any domain where outcomes are comparable under fixed quality and service constraints: cost per resolved ticket against cost per resolved ticket, never cost per token as a universal solvent. The field has flagged its own version of this bracket, the move from cost per token to cost per outcome, among the same founding questions, so the restriction is an insider acknowledgment rather than an external caveat. In operational form, this paper's proposal accordingly replaces intelligence per watt with quality-adjusted outcome per joule, and per watt of provisioned capacity where capacity binds; Section 7 gives the denominator its own contract.

With units fixed, the structure of the measurement problem is three ledgers. The financial ledger, in dollars, is driven by hardware or cloud pricing, energy price, labor, networking, and utilization; it is the ledger the token bill partially meters and the one every existing budget process reads. The energy and capacity ledger, in kilowatt-hours and megawatts, is driven by hardware efficiency, utilization, workload volume, and facility design; it is independently material even on a perfectly clean grid, because power constrains before carbon does, which is the field's own claim about the binding constraint, restated with units attached. A firm that cannot get a grid interconnection has a problem of energy and capacity that no emissions factor and no dollar figure describes. The emissions ledger, in CO2-equivalent, holds the physical accounts: the attributional operational account displayed above; an embodied account, arriving amortized inside hardware prices that do not disclose it; and, where the question is causal, consequential estimates, a different measurement class, carried separately and flagged as such. Contractual instruments, certificates, power purchase agreements, allowances, credits, do not belong inside a ledger denominated in CO2-equivalent at all, precisely because they are claims held against the ledger rather than measurements of it. They live in an instrument register beside the ledgers, each entry carrying its type, vintage, volume, matching basis, eligibility, and retirement status. A third kind of object completes the structure: derived reporting results, market-based inventory figures among them, which are the outputs of a named accounting rule applied to physical accounts and register together, reproducible from their inputs and never primitives. The emissions ledger is the one whose conversion into either of the others is the subject of this paper, and the separation is why Section 4 can give four defensible descriptions of one consumption: three read accounts, and one is a result derived from a register position under a named rule.

It is tempting to compress the three into one activity base times different rates, the same usage carrying a dollar rate and a carbon rate. That compression is a serviceable approximation and Section 5 will use it as such, but the data model underneath is not that clean. Embodied carbon has no dollar line of its own; it arrives amortized inside a hardware price that does not disclose it. Several dollar lines, software licensing, labor, network egress, have no energy content at all. And the observability differs by position in the stack: an API customer observes a token bill in dollars while the physical energy behind it is invisible to them, which means the financial ledger can be fully populated at a boundary where the energy ledger is empty. Three ledgers, then, not one number seen three ways: distinct drivers, distinct visibility, distinct units. The separation is not novel bookkeeping. Environmental management accounting has maintained physical and monetary accounts side by side for two decades under the name carbon management accounting (Stechemesser and Guenther 2012), and emissions measurement frameworks at workload level already run the physical ledger for compute specifically (Hanafy et al. 2023). What this paper adds to that ancestry is capacity made independently visible within the energy and capacity ledger, and the argument for why the separations must survive the pressure to collapse them.

Assembled, the measurement object is a vector. For a workload with outcome Y:

$$M(Y) \;=\; \left(\frac{C}{Y},\ \frac{K}{Y},\ P_\tau,\ \frac{G_A}{Y},\ \frac{G_E}{Y},\ \frac{G_Q}{Y}\right)$$

with C total financial cost, its decomposition disclosed, K energy, P_τ power or capacity exposure over a defined interval τ, G_A attributional operational emissions, G_E embodied emissions, and G_Q a consequential estimate where the question requires one; the instrument register R travels beside M, not inside it. Measurement produces this vector. The pressure on it has a precise form: a single score requires a conversion factor between units, and the requirement is where the trouble lives.

3. No canonical exchange rate

Any combined metric is not a measurement derived from M; it is a decision functional applied to M:

$$D(M, R;\ \lambda,\ b,\ p) \ = \ \frac{C + \lambda\, A_b}{Y}$$

where R is the instrument register, b names the emissions basis, $A_b = A_b(M, R)$ is the quantity that basis produces, a physical account read directly or a derived result computed under a named rule, λ is a carbon value in dollars per tonne, and p is the decision purpose. Admissibility is part of the definition: the pair (λ, b) must belong to $\mathcal{A}(p)$, the set of valuation and basis combinations valid for the purpose at hand, and the survey below shows why most cross pairings are category errors. C is total financial cost, carried with a disclosed decomposition, energy, infrastructure, labor, policy payments, claim purchases, with the register associating carbon related expenditures to their entries rather than every register expenditure adding mechanically. The notation forces four hidden selections into view: which emissions quantity, produced by which rule, valued at which price, for which purpose. Measurement produces the vector; governance produces the scalar. The question is not whether carbon can ever be priced; it is priced constantly, in many places. The question is what authorizes the choice of λ for the decision at hand. Double counting is a semantic rule, not an arithmetic one: a cash flow for carbon and a monetized emissions quantity may legitimately coexist when they serve different purposes, an allowance cost paid for compliance alongside a λ set at damages applied for appraisal, and duplication occurs only when the same obligation or valuation enters twice, which is exactly what declaring p makes checkable. The disclosed decomposition earns its keep on the hard cases: a power purchase agreement that is procurement and instrument at once, a bundled certificate with no separable price, a carbon tax that is carbon expenditure but no register entry. And one boundary rule completes it: a market price may enter as an observed cash cost inside C, procurement, or as a scenario valuation λ applied to A_b, appraisal, but never both for the same obligation. Credit purchase is the procurement case in full, eligible instrument volume times price, a cash line in C's decomposition rather than a monetization of any emissions account; allowance expenditure already inside C is the same case, and the compliance row's λ must not price it again. Surveying the candidates shows why no general authorization exists.

Instrument	What it represents, and where it applies	What it cannot establish
Compliance allowance price (EU ETS ∼ 65 average in 2024; ∼ 80 in mid-2026)	The clearing cost of a permit to emit inside a capped system; applies to covered installations in covered jurisdictions, roughly a quarter of global emissions	The damage a tonne causes; the cost of a tonne outside the cap; a value for a firm no compliance system covers
Carbon tax rate	A legislated levy per tonne in a taxing jurisdiction; applies where the statute applies	Anything beyond the jurisdiction's political settlement; rates vary widely and change with governments
Voluntary credit price ($6.34 blended average in 2024; $4.64 for reductions against $17.28 for removals)	The transaction price of a claim instrument in an uncapped market; applies to the claim purchased	Physical equivalence to a tonne of one's own emissions; the integrity evidence below is the reason
Social cost of carbon ($190 per tonne at a 2 percent discount rate in the 2023 federal estimate; a $140 to $380 family for 2030 across discount rates)	An estimate of the marginal damage of a tonne, for policy appraisal; applies to benefit-cost analysis of public action	A single authoritative figure; the estimate is parameterized by contested normative choices and its official selection is politically unstable
Internal carbon price (roughly two thousand firms; median near $49, set across a wide range)	A planning value a firm imposes on itself to prepare decisions for future regimes; applies inside the firm that set it	A market fact; it is a governance choice, and it is frequently set above every observable market price on purpose
Marginal abatement cost	What the next avoided tonne costs a specific actor by a specific action; applies to that actor's abatement ordering	A price for tonnes in general; it is a schedule, not a scalar, and it differs across actors by orders of magnitude

Sources for the table's levels: compliance and coverage per World Bank (2025) and ICE market data as of the dates shown; voluntary prices and the removals split per Ecosystem Marketplace (2025); the social cost family per EPA (2023), the range per Congressional Research Service (2025), and the discount sensitivity per Tol (2023); internal prices per CDP (2023).

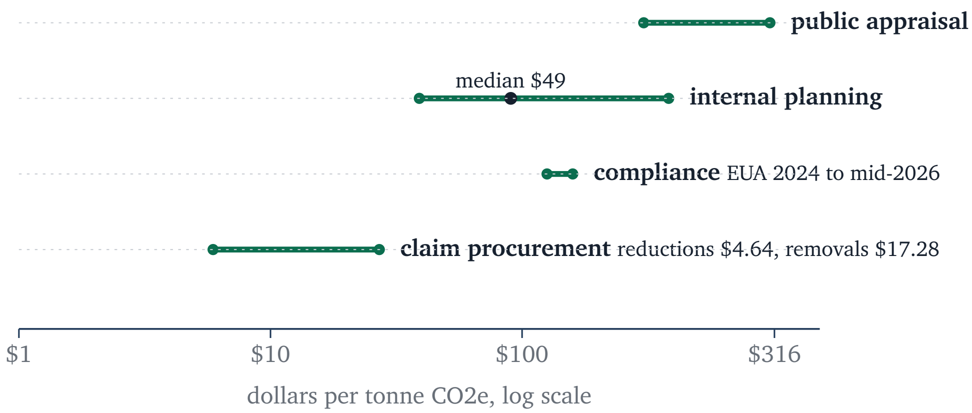

Figure 2. Four purposes, four price ranges, one axis. The vertical axis carries no quantity: each band is a different question, and the horizontal spread within and between them is the paper's central observation, one to two orders of magnitude between the cheapest claim and the highest appraisal value. The bands are deliberately not stacked as a single distribution, because these are not competing estimates of one underlying number. A price is admissible only for its own row's purpose, against that row's emissions basis; reading across the figure horizontally, as though a claim price and an appraisal value were substitutes, is the category error Section 3 describes. Levels as reported in Section 3, dated 2024 to mid-2026 and shown in their native currencies converted only for placement.

The instruments disagree by one to two orders of magnitude, and the disagreement is the point. These numbers do not differ because one market is malfunctioning while the others work. They differ because they answer different questions: what does a permit cost, what did the legislature set, what does a claim trade for, what does a tonne do to the world, what does this firm assume, what does abatement cost this actor. Each value is meaningful within its purpose and undefined outside it. That is purpose-specificity, and it must not be mistaken for relativism: the third column is the discipline. A voluntary credit price is a perfectly good answer to what the claim costs and an indefensible answer to what the emission did; a social cost estimate is built for public appraisal and is not a compliance obligation anywhere. Local to a purpose does not mean equally defensible for a given purpose.

The voluntary market's integrity evidence belongs here, repurposed. It is often read as proof that the market is fake; the reading this paper needs is narrower and better supported: a transaction price cannot serve as a factor for physical equivalence. Independent evaluations of the largest project category find that roughly seven to twenty-five percent of credited reductions are real and additional (West et al. 2023; Probst et al. 2024), and a 2026 synthesis across forty-four projects and six independent evaluation methods finds aggregate overcrediting near elevenfold, alongside the finding that four in five projects did reduce deforestation (Swinfield et al. 2026), integrity failure and real activity coexisting. A study of the twenty largest corporate buyers finds eighty-seven percent of their holdings rated high risk, selected, the authors conclude, substantially on price (Trencher et al. 2024). And the market itself now prices the integrity difference: removals transacted at a premium of roughly three hundred and eighty percent over reductions in 2024, up from about two hundred and forty-five percent the year before (Ecosystem Marketplace 2025). Integrity heterogeneity is a leading driver of this price structure, among others; the causal claim need go no further. The pattern has a name in the market design literature, a lemons market, in which unobservable quality drags the traded price toward the worst supply (Akerlof 1970); the framing is established prior art for the voluntary market (Haya et al., Berkeley Carbon Trading Project), and this paper leans on it rather than deriving it afresh.

The social cost of carbon fails as a universal λ for a different reason: not integrity but authority. It is a family of estimates parameterized by choices no measurement can settle, the discount rate above all, jurisdictional scope and damage functions besides, and the sensitivity is large: moving the discount rate from two percent to one and a half takes the 2020-emissions figure from \$190 to roughly \$340 (EPA 2023; Tol 2023). There is no authoritative selection an operator can defer to, and the one instrument that made a federal selection was dismantled through 2025: by executive order, agency revisit, a directive across government against factoring climate damage into regulation except where statute requires, and a departmental report arguing the value is too high (sequence per Harvard Environmental and Energy Law Program 2026). The dismantling is evidence of policy instability in the selection of λ, not evidence that the damages ceased to be meaningful; the domestic only framings under discussion imply single digit dollar values, which would place the highest of the three public prices on a trajectory toward the lowest. For present purposes the direction of travel matters less than the fact of it: the anchor an operator might have deferred to is itself in motion.

The strongest adjacent analysis reaches this paper's terrain from the welfare side, and the relationship states cleanly in λ vocabulary. That analysis argues that AI production buys a regulated input, electricity, at prices below marginal social cost, and sells an unregulated output, tokens, at prices that embed no externality, an implicit subsidy flowing from ratepayers and the environment to compute (Litowitz, Polson, and Sokolov 2026). On that analysis's showing, the transfer exists whether or not anyone can name its dollar size. But the remedy, internalizing the externality, requires selecting a λ, and the selection is exactly the purpose dependent step just shown to be unsolved in general. This paper concedes the diagnosis in full and takes only the remedy's missing step as its subject. One adjacent term should be retired before it confuses: carbon currency exists in the literature as a proposal to make carbon itself a global monetary instrument (Liu, Chen, and Xiao 2022), which is close to the opposite of the observation made here, that no single conversion governs the firm's decision.

The admissibility set can be made operational in one compact table, illustrative rather than exhaustive:

Decision purpose	Relevant quantity or accounting basis	Applicable price or value
Compliance budgeting	Statutorily covered emissions	Allowance price or tax rate
Public benefit-cost analysis	Consequential or policy defined emissions	Social cost of carbon
Internal transition planning	Disclosed physical or regulatory basis	Internal scenario price
Abatement prioritization	Avoided emissions of each intervention	Marginal abatement cost
Voluntary claim procurement	Eligible instrument volume	Credit purchase price

Read against the earlier table, this is the same content inverted: there, each value's third column said what it cannot establish; here, each purpose's row says what it can. The last row is the procurement case: its price enters as cash in C, and the compatibility discipline governs which claims the purchase can support rather than which λ applies. The principle the two tables share is the section's conclusion in positive form: the emissions basis, the carbon value, and the decision purpose must be mutually compatible.

The design principle follows, and Sections 5 and 7 cash it. Any metric that silently selects one of these valuations has made a governance decision and filed it as arithmetic. A blended score with an undisclosed λ presents a compliance price, or a claim price, or a planning assumption, as arithmetic. A discipline writing its standards now has a choice its predecessors lacked: it can make the selection explicit instead. In the venue's own terms: carbon valuation is a policy or governance input, not a universal measurement constant.

Purpose-specificity would remain an abstraction if the purposes never met. In electricity emissions accounting they meet on the same kilowatt-hour, and a live standards process shows what happens.

4. Scope 2 as the case study in nonequivalence

Take one workload, in one facility, for one month, and describe its electricity four ways. First, an electricity bill: so many dollars for so many kilowatt-hours. Second, an annual contractual claim: the operator holds renewable energy certificates or a power purchase agreement covering one hundred percent of annual consumption, and reports the consumption as clean under the market-based method. Third, an hourly, locational emissions profile: the same consumption, matched hour by hour against the grid that serves the facility, showing which kilowatt-hours coincided with clean generation and which drew on the evening gas fleet. Fourth, a consequential estimate: what generation, somewhere, ramped because this load existed. All four are defensible. All four are answers to different questions, the financial, the contractual, the physical, and the causal; in Section 2's architecture, the first reads the financial ledger, the third and fourth read two different emissions accounts, and the second is a derived reporting result, the market-based rule applied to a register position. No arithmetic reconciles them without first choosing which question wins.

The taxonomy is the live subject of the standards process governing corporate electricity accounting. The market-based method rests today on annual certificate matching. The revision consultation that closed on 31 January 2026, with over a thousand submissions from thirty-nine countries, proposes hourly matching and deliverability requirements for market-based claims, alongside a cap on claims to standard supply, load profile approximations, exemption thresholds, a legacy clause for existing contracts, and a multiyear phased implementation; a second consultation follows in 2026, with the final standard expected around 2027 and implementation around 2028 (GHG Protocol 2025, 2026). If adopted in substantially the proposed form, some annual instruments could cease to qualify for certain market-based claims. Nothing here asserts the outcome; a consultation is not a standard. What the process demonstrates is already sufficient: the standard setter itself treats the annual contractual and hourly physical descriptions as different reporting objects requiring different rules.

The category distinction underneath does real work. A certificate can be a fully valid attributional claim to an energy attribute without being evidence of hourly matching or of additional generation; in this paper's terms it is a register entry from which a derived result is computed under a named rule, not an emissions account, and the standards work exists precisely to keep physical accounts, instrument claims, and derived results distinct, with certificates and offsets best understood as related integrity pressures on different register entries. The gap the revision addresses is illustrative rather than scandalous: a facility on a solar purchase agreement can report full annual coverage while a minority of its consumption matches clean generation hour by hour, solar peaking in the Texas afternoon while the load runs through the Virginia midnight, and independent analysis has questioned how much real effect the current annual regime delivers (WattTime 2025). The proposal's own feasibility analysis makes the incidence plain: under the thresholds being considered, most reporting companies would be exempt from hourly matching while the vast majority of grid load would remain covered (GHG Protocol 2025), which is to say the tightening is aimed at exactly the class of large, concentrated loads that AI facilities are.

The AI bite comes from the physics of the grid meeting the geography of the buildout, and the figures carry estimate flags. Published analyses put grid carbon intensity varying by roughly a factor of two across the hours of a day and by tens of times across regions, on the order of forty-fold between the cleanest and dirtiest major grids (Electricity Maps data; Hanafy et al. 2023). AI load, meanwhile, is concentrating in a handful of transmission regions where the spread between the contractual claim and the hourly profile runs widest (NERC 2025; Goldman Sachs Research 2026b). An industry whose consumption is large, growing, geographically concentrated, and temporally flexible is the industry on which the granularity question lands first, and the one with the most to gain from measuring at the granularity the question demands.

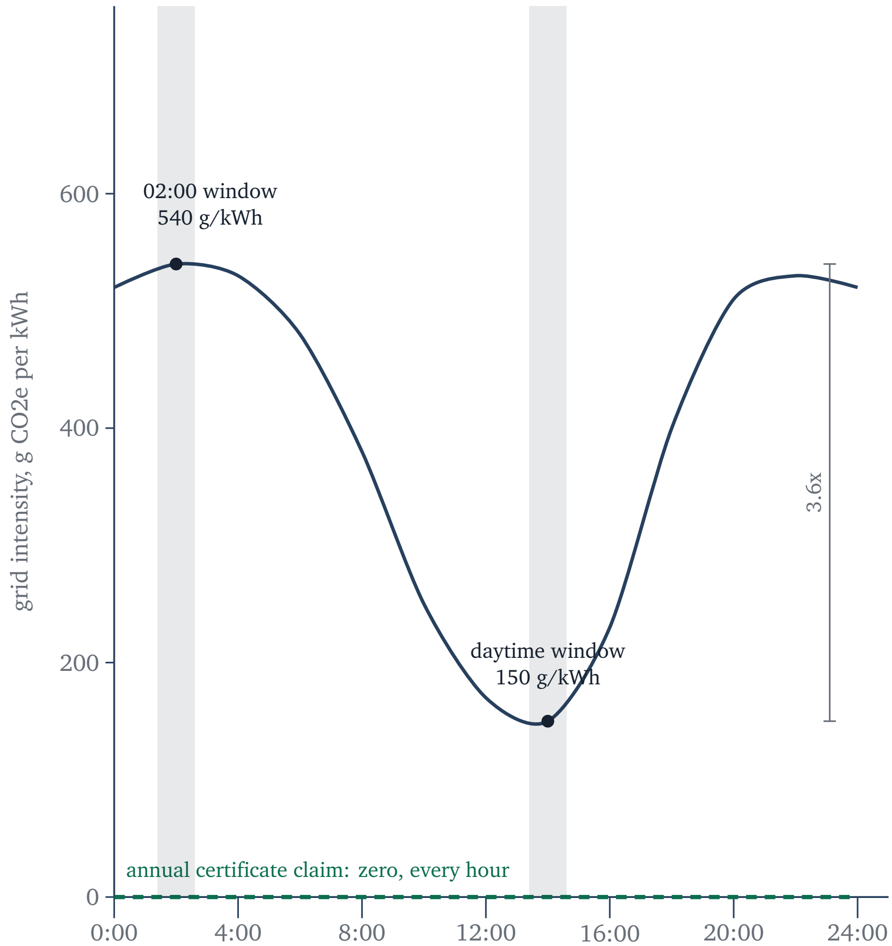

Figure 3. Two reporting objects, one day of consumption. The curve is a day of grid carbon intensity on a solar-heavy system; the dashed line is what an annually matched certificate reports for every hour of that same day. Identical work carries a factor of several difference depending on the hour it runs, while the contractual account reports one number throughout. Neither line is wrong. They answer different questions, which is why the revision described in this section treats them as different reporting objects rather than as competing estimates of one quantity, and why the worked case in Section 5 can rank the same options differently on each basis. Curve shape illustrative; the marked hours correspond to the worked example's daytime and 02:00 windows.

The landing is the thesis made operational. What Scope 2's revision exposes is not fraud but divergence: financial, contractual, and physical accounts of identical consumption legitimately differ, because they answer different questions about it. That is what a vector-valued object looks like in the wild. A discipline standardizing measurement now should draw the corresponding conclusion: standardize the components, their boundaries, and their disclosure, rather than a collapse that would freeze one question's answer into every question's number.

If the accounts legitimately diverge, decisions still get made. Three canonical architectures carry nearly all of that practice, and together with the ledgers they assemble into something the discipline can adopt.

5. The decision architecture, and the stack it assembles

None of what follows is new machinery. Every element is in use somewhere today; the contribution is the organization, the application to compute, and the assembly at the end. Three architectures make the carbon decision explicit without pretending a universal price exists.

The first works from a constraint. Minimize financial cost subject to an emissions ceiling, alongside the quality, latency, reliability, and capacity constraints every serious deployment already carries. This is the natural form wherever a budget or target exists, a corporate commitment, a jurisdictional cap, a customer requirement, and it has a property the other two architectures lack: in a continuous convex formulation, the shadow price of the binding emissions constraint is the decision's own local λ, produced by the optimization rather than assumed into it; in discrete deployment choices like this paper's example, the analogous object is a finite-difference value, or a range of values consistent with the selected option, rather than a unique multiplier. The concept has a long lineage in environmental economics, where the shadow price of a carbon budget is the scarcity value of the constraint in its own units (Pezzey 2019); the architecture lets each deployment discover its own, in whichever form the problem admits.

The second works from the Pareto frontier. Where no ceiling has been set, surface the options no alternative beats on both cost and emissions, and present the frontier. This defers the λ choice to the decision maker while eliminating every option dominated on both cost and emissions, retaining all efficient alternatives, including unsupported ones no linear weighted sum would select, and it makes the trade-off's shape visible: sometimes the frontier is steep and a small cost concession buys a large emissions reduction, sometimes it is flat and the concession buys little. The mechanism's home is the carbon-aware scheduling literature, which has been computing cost and emissions frontiers for cloud workloads for years (Wiesner et al. 2021; Hanafy et al. 2023); the addition here is the governance framing, the frontier as the artifact a decision maker signs rather than an engineering curiosity.

The third works from scenario prices. Evaluate the decision under several values of λ, from zero through a high policy value, and ask not which price is true but whether the decision changes across the plausible range. If the same option wins everywhere, the carbon question did not bear on the decision and the analysis says so. If the winner flips, the flip point is the number the governance conversation is really about. This architecture has a documented constituency that already practices it in spirit: on the order of two thousand firms maintain internal carbon prices, set across a wide range and frequently above every observable market price, explicitly to prepare decisions for regimes that do not yet bind them (CDP 2023). That practice is evidence for this paper's premise from the sophisticated end of the market: participants with the most at stake already decline to let the market's λ make their decisions.

The line all three architectures build through is this: separate ledgers do not eliminate judgment; they make judgment visible. The constraint architecture exposes the scarcity value of a binding ceiling where the problem admits one; the frontier preserves the full efficient option set, unsupported options included, before any weighting is applied; and the scenario architecture shows where a disclosed valuation changes the choice. What none of them does is let the judgment disappear into a blended number. The three are canonical rather than exhaustive, lexicographic rules and multicriteria methods exist, but each of these corresponds to a governance posture firms already run, which is what a standard can bind to.

The architectures earn their keep on a distinction that ordinary cost optimization does not make. Quantity reductions, eliminating idle, raising utilization, spending fewer tokens per outcome, using smaller or specialized models, caching, batching, tend to move all three ledgers together: less work is less money, less energy, less carbon (this is where Section 2's approximation, the same usage carrying two rates, is good enough). These are largely captured by ordinary dollar optimization, and the discipline's existing playbook already hunts them. Composition changes are different: shifting work across time and regions, choosing hardware generations, structuring procurement and power contracts, siting storage and onsite generation, preserving portability. These move the ledgers independently and sometimes oppositely, a shift off peak that cuts carbon at flat cost, a region move that cuts cost at higher carbon, and the dollar ledger may move little or not at all when they happen, which is precisely why dollar optimization does not reliably surface them. The composition class is where the emissions ledger pays for its keep, and where a decision architecture, rather than a reflex, is required.

The distinction, and the architectures, can be watched working in one compact case. Take one workload, a million documents processed per month at a fixed quality bar, and four ways to run it, with figures per thousand outcomes; each batch of a thousand is processed inside a window of one hour, on hardware whose instantaneous draw is bursty around an 8 kW average with a 45 kW provisioned peak. Four power quantities must be held apart: the workload's instantaneous draw, its provisioned capacity (unchanged by scheduling), its contribution to the facility's coincident peak window (a defined daytime hour), and its contribution to the grid system peak. The table's last column reports the third. The values are constructed solely to demonstrate the mechanics and should not be interpreted as estimates of any particular workload.

Option	Cost	Energy (kWh)	Physical emissions (kg, hourly locational)	Market-based result under stated rules (kg)	Contribution to facility coincident peak (kW)
A: cheap grid region, daytime	$10.00	8.0	4.80	4.80	45
B: option A plus annual renewable certificates	$10.80	8.0	4.80	~0	45
C: cleaner region, daytime	$10.50	8.2	1.23	1.23	45
A : option A shifted to 02:00	$9.70	8.1	5.35	5.35	0

Register note: option B's entry is denominated in energy attributes, certificates retired covering 8.0 kWh equivalent per thousand outcomes on annual matching with stated vintage and eligibility; the near zero kilogram figure in the derived column is the inventory result of applying the market-based rule to that position, not the instrument itself.

The table supports three readings, taken in turn, with the feasible set stated at each door.

First reading: accounting basis, capacity aside, A held back. Read A, B, and C against each other, and the paper's claims stop being abstract. Lower energy does not mean lower emissions: A uses less energy than C and emits nearly four times as much, because the factor differs. Lower cost does not mean lower anything else: A wins that comparison's dollar column and loses its physical one.

And the accounting bases redraw the frontier itself. On the market-based basis nothing dominates: B reports the lowest claim at the highest cost, C the middle of both, A the cheapest and highest, three points on one frontier, and selection still requires a ceiling or a value, with the flip between B and C sitting near $\lambda^*_{BC} \approx \$244$ per tonne on that basis (the thirty cent cost gap over the 1.23 kg claim gap), arithmetic any basis supports, reported as a mathematical sensitivity of the derived result rather than a recommended valuation; which purposes admit monetizing a derived reporting result at all is exactly the compatibility question Section 3 poses. On the hourly physical basis the frontier collapses to A and C, because B's certificates change its register position and hence its derived result, not its physical account, leaving it dominated, the same emissions as A at higher cost. The basis choice does not merely reorder options; it changes which options are candidates at all.

A ceiling decides without any price: an emissions budget of 2.0 kg per thousand outcomes leaves C as the only feasible option. **Second reading: scalarization against the frontier, capacity still aside, A entering.** Where a scalar is insisted on, the flip price on the physical basis is computable:

$$\lambda^* \;=\; \frac{C_C - C_A}{E_A - E_C} \;=\; \frac{\$0.50}{3.57\ \text{kg}} \;\approx\; \$140 \text{ per tonne}$$

Below roughly $140 per tonne, the blended score selects A; above it, C. Reading that threshold obeys Section 3's condition. The purpose is internal transition planning against the hourly attributional basis, and for that purpose management selects and discloses a scenario range: observable compliance prices and policy appraisal values may inform the governance discussion, but they are not interchangeable estimates of one underlying price, a credit price is procurement inside C and never a λ here at all, and none of them is the range. Run against a disclosed illustrative range of $0 to $250 per tonne, the threshold sits inside it, so the planning decision turns on where governance draws its range, not on the measurements: the compatibility condition, demonstrated. Restore A , and composition becomes concrete and uncomfortable: the same batch shifted to 02:00 in the same region, where the hourly attributional intensity is higher once solar leaves the mix. Cost falls, the contribution to the facility's coincident peak collapses toward zero while the workload's own overnight draw is unchanged, and physical emissions rise; one scheduling decision moves the ledgers in opposite directions. And A resets the scalar landscape it was not part of: with it present, the blended score selects A below roughly $194 per tonne and C above, while A, though Pareto-efficient on the physical frontier, is unsupported, sitting on a nonconvex stretch of the discrete frontier where no single λ ever selects it. Stated generally: linear scalarization by a carbon price selects only supported efficient options, and in a discrete, nonconvex option set it can miss Pareto-efficient options entirely. That is a substantive argument for the frontier architecture of this section, not a curiosity of the example: Pareto analysis retains efficient discrete options that no linear carbon price will ever surface.

Third reading: feasibility. It comes last on the page because in practice it comes first. Feasibility turns on the facility, hour by hour, because an interconnection limit applies to total instantaneous load in every hour, not only the peak one. Give each facility a baseline:

Period and facility	Facility baseline (kW)	Workload (kW)	Total (kW)	Interconnection limit (kW)
Daytime, facility hosting A and B	970	45	1,015	1,000
Daytime, facility hosting C	940	45	985	1,000
02:00, same facility as A	700	45	745	1,000

At the first facility the daytime headroom is 30 kW against a 45 kW need, so A and B are infeasible as sited; C's facility clears with 15 kW to spare, and the 02:00 hour clears with 255. Feasibility makes the first cut, pruning the set to C and A before any price or ceiling is consulted, the workload's own 45 kW simply moved to where headroom exists; the \$194 threshold then prices the survivors, and where the 2.0 kg ceiling binds as well, C is the choice outright, A the first option excluded. That is the result the architecture exists to govern: decision rankings are unstable across accounting bases, decision rules, and even the option set, under conditions the disclosures make visible.

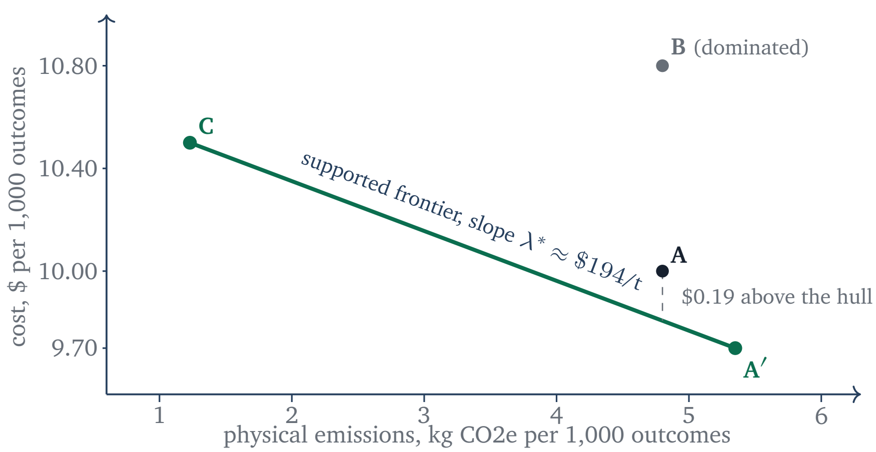

Figure 4. The frontier, and the option no price selects. Each point is one deployment option from the worked case, read on the hourly physical basis. B is dominated: same emissions as A, higher cost. A, C, and A′ are all Pareto-efficient, yet the supported frontier, the lower hull that linear scalarization can reach, runs from A′ straight to C. A lies \$0.19 above that chord, so no carbon price whatsoever selects it: a weighted sum sweeping λ from zero upward jumps from A′ to C at roughly \$194 per tonne and never stops at A. Frontier analysis retains the option; scalarization cannot see it.

Assembled, the pieces are a stack, and naming it is this paper's constructive close.

Layer	Measure and unit	Disclosure the layer requires
Output	The quality-adjusted outcome Y, specific to the task	The outcome definition and the quality and service constraints held fixed
Financial	Total cost of ownership per outcome, in dollars	Cost boundary and decomposition: what is inside (hardware, energy, labor, licenses, policy payments, claim purchases), what is amortized, and which expenditures associate to register entries
Energy and capacity	Kilowatt-hours or joules per outcome; useful output per watt of provisioned capacity where capacity binds	Measurement boundary and whether figures are metered, modeled, or reported by a vendor
Emissions	CO2-equivalent per outcome	Temporal and geographic boundary; attributional or consequential basis; operational and embodied split; embodied allocation assumptions (asset life, utilization, capacity share)
Instrument register	Certificates, power purchase agreements, allowances, credits held against the book	Instrument type, vintage, volume, matching basis (annual or hourly), eligibility standard, retirement status
Derived reporting results	Inventory figures computed under named standards, market-based Scope 2 among them	The standard and rule version applied; the accounts and register entries consumed; reproducibility of the computation
Decision	The rule in force: constraint, frontier, or scenario prices; or a disclosed organizational shadow price	The rule itself, and the λ or ceiling it commits to, stated

The stack is not a proposal built from nothing. The output layer is Section 2's bracket made operational. The financial and energy layers are the field's two founding metrics, given boundaries. The emissions layer is Section 4's physical components with their disclosures attached, and the register and derived results rows are Section 2's separation kept structural: claims beside measurements, reported figures reproducible from both, never primitives. The decision layer is this section. Every layer is already in use severally, in the field's own materials, in carbon management accounting, in emissions measurement at workload level, in the scheduling literature, in internal pricing practice. The stack formalizes the architecture implicit in that practice; it is that practice made composable.

The stack handles a decision at a time. The harder question is posture, because the rules of the emissions layer are visibly in motion, and the stack is also what makes the motion navigable.

6. Operating while the rules move

Begin with what each layer of the stack reveals, in the discipline's own vocabulary. The financial and energy layers, read together, expose idle capacity and utilization failure, which is the discipline's visibility and unit economics work. The emissions layer, carried at hourly and regional granularity, exposes the composition opportunities of Section 5 that the dollar ledger cannot see, which is optimization and workload placement. The decision layer's scenario prices expose sensitivity to policy futures before those futures bind, which is accountability. And portability and observability, the properties that keep the stack populatable as standards change, preserve the capacity to respond, which is optionality. The operator running all layers is the mature three-ledger operator the discipline's own principles describe, doing with three ledgers what it already does with one.

The motion is documented rather than predicted. The Scope 2 revision runs on its published timeline, second consultation in 2026, final standard expected around 2027, implementation around 2028 (GHG Protocol 2026). The removals premium in the voluntary market widened from roughly two hundred and forty-five percent to nearly four hundred percent in a single year as claim standards converged on durability (Ecosystem Marketplace 2025). Compliance coverage stands near a quarter of global emissions and is expanding through border adjustment, with the reference price for the first border adjusted quarters now published (World Bank 2025; European Commission 2026). Three mechanisms, each contingent, at present pointing the same way: the rules governing the emissions ledger are tightening on published schedules, and published schedules can slip.

Two boundary markers keep this section on its own ground. The established machinery for carbon under uncertainty, timing options and internal prices, presumes a single governing price process that wanders; the condition here is different, several prices governing different purposes with the applicable one undetermined, and that condition is answered below by classification rather than optimization. And the categories that follow descend from the robust-decision-making tradition, which asks what performs across futures rather than what is optimal in one (Lempert, Popper, and Bankes 2003); the debt is acknowledged and the machinery kept light.

The classification has three classes. No-regret actions perform across nearly all regimes: measurement infrastructure, a workload inventory, idle elimination, efficiency investments with clear payback, observability, and the avoidance of gratuitous lock in. These actions can often be justified on the financial and energy ledgers alone, while also supplying much of the activity data the emissions ledger needs. Option-preserving actions are valuable for the response capacity they hold open rather than the regime they assume: portability across regions and providers, temporal flexibility in scheduling, contractual access to cleaner supply, an executable abatement option set backed by committed liquidity, named exactly so because an option without committed liquidity is a slide rather than a capability, and supplier diversification. Forecast-dependent actions ride on one future: large prepaid positions in specific instruments, infrastructure justified by a single regulatory forecast, and reliance on the permanent validity of a particular claim standard. Validity risk, the substance of that last item, is a spectrum rather than a cliff: an instrument can face future ineligibility under a revised standard, reputational reassessment while remaining formally valid, contractual replacement obligations, or, rarest, formal restatement of past claims; transition and legacy provisions are common in standards processes, including the one in Section 4, and the evidenced weight today sits at the reputational end of that spectrum. The classes support one conclusion and no more: under deep uncertainty about the correction mechanisms, no-regret and option-preserving investments typically justify themselves without a view on the regime, and forecast-dependent commitments are bets, to be sized as bets.

How the postures fare across futures can be stress tested qualitatively. The matrix below is exactly that, an illustrative stress test of four decision rules against four regimes, not a payoff calculation.

Decision rule	No material change	Granular accounting adopted broadly	Binding price or cap	Claim-standard tightening, legacy uncertain
Financial optimization alone	Low exposure; lowest measurement cost	Retrofit measurement under deadline; moderate adaptation cost	Exposed to repricing of unpriced carbon; high transition cost	Claims never made, so little restatement risk; reputational lag
Compliance-minimum	Low exposure	Moderate adaptation cost where thresholds bind	Adapts by construction, at prevailing prices	Exposure where minimum relied on legacy instruments
Carbon-first precommitment	Overcommitted; visible under-performance against peers	Well positioned if instruments chosen match the granular rules	Positioned, with stranded position risk in the wrong instruments	Stranded cost risk concentrated in early instrument choices
Mature three-ledger governance	Incurs measurement and governance overhead; may underperform cost optimization alone where no carbon linked regime or decision benefit materializes	Lower adaptation cost where the granularity already exists	Prepared where scenario prices bounded the exposure in advance	Prepared where the validity spectrum was priced as a spectrum

The interim fairness the matrix owes is in its first column and should be said plainly: the three-ledger rule incurs real measurement and governance overhead in the regime where nothing changes and may underperform optimization on cost alone where that overhead produces no offsetting operational or decision benefit, and any account that hides that cost is selling rather than analyzing. What the remaining columns record is what the overhead buys. The matrix is a reading aid, not a derivation; readers should weigh the regimes by their own probabilities, which the framework deliberately does not supply.

A closing inference, marked as inference. Organizations' preferred measurement and policy rules may reflect their installed capabilities as well as their stated principles: an operator with hourly telemetry gains from hourly standards, an operator with a legacy certificate book gains from legacy clauses, and neither preference is insincere for being convenient.

It does mean that debates over measurement standards are also, quietly, debates over who holds advantage under the next rules, which is one more reason the selection of rules should happen in the open, in a standards process, rather than inside a blended metric.

The implementation implication follows from the model with no further assumptions: because accounting standards and applicable valuations can change, the architecture preserves raw measurements, register entries, rule versions, and decision assumptions, so any reported result can be recomputed under new rules without reconstructing the underlying workload data. That property, recomputability, is what the postures above protect.

What remains is what the answer changes, for the measurement, the management, the capital, and the discipline that asked.

7. The Minimum AI Compute Impact Reporting Profile

The deliverable is a set of interoperability rules for the measures the industry already uses, offered in the form standards guidance takes and continuous with the standardization the field has already begun:

1. Never report TCO per output without identifying the output and the quality constraint it was held to.
2. Never report energy efficiency as a carbon measure; the conversion requires an emissions factor, and the factor requires a time, a place, and a basis.
3. Never report monetized carbon without disclosing the selected value and its purpose, and never without the underlying physical quantity alongside.
4. Distinguish operational from embodied emissions, and energy from power, wherever either pair appears.
5. Define the temporal and geographic boundary of every emissions figure, state whether its basis is attributional or consequential, and report contractual instruments in a register beside the physical accounts, never netted silently into them, labeling any market-based or other derived figure with the rule that produced it.
6. Distinguish optimization that reduces computational demand from optimization that changes energy composition; the first moves the ledgers together, the second does not.
7. Permit multiple decision methods; standardize the disclosure of the method, not a universal carbon price.
8. Never apply a carbon value to an emissions quantity without naming the decision purpose; the basis, the value, and the purpose must be mutually compatible.

These are the same standardization instinct that has already extended cost specifications to token economics (FinOps Foundation 2024; Linux Foundation 2026), applied one layer further.

The denominator needs the same rigor as the numerators, because Y is the profile's central comparability assumption, and because the field's first framework already relies on it: Big-T compares complexity classes at a stated quality floor, quality the precondition of the comparison rather than a term in it (Neff 2026). The contract gives that floor its fields, versions, and coverage rule, and the two artifacts compose, Big-T describing how consumption grows, this profile stating what each unit of it carries in cost, energy, and emissions. An Outcome Contract specifies it:

Outcome contract field	Requirement
Unit of outcome	Precisely defined
Evaluation population	Dataset, traffic cohort, or sampling rule
Success criterion	Metric and threshold
Failure treatment	Retries, abstentions, unresolved cases
Human contribution	Review, correction, and escalation boundary
Service constraints	Latency, reliability, availability
Statistical basis	Sample size and confidence or error range
Case flow and coverage	Eligible, attempted, succeeded, abstained, failed, escalated; named rates: attempt rate, successful automation coverage, conditional success rate
Version	Definition, benchmark, effective date

And the rules compile into an implementable artifact, a Minimum AI Compute Impact Reporting Profile, one record per workload:

Field	Requirement
Outcome definition	Required; an Outcome Contract, versioned
Case flow and coverage	Required: counts per contract row, successful automation coverage beside every per outcome figure, and impact per eligible case as a required companion figure wherever that coverage is below 100 percent
Quality threshold and service constraints	Required
Financial boundary and decomposition	Required enumeration
Energy per outcome	Required, or explicitly unavailable with reason
Power or capacity metric	Required where capacity is material; interval stated
System boundary	Required: automation subsystem only, or end-to-end service including human review
Geographic and temporal boundary	Required; resolution stated
Emissions method	Required; named method and version
Data provenance	Required per figure: metered, modeled, allocated, or reported by a vendor
Embodied allocation	Required where embodied is reported: asset life, utilization, capacity share
Instrument records	Separate, itemized register entries
Derived results	Named rule and version; inputs identified for recomputation
Decision purpose	Required when any monetization is reported
Carbon value	Value, currency, date, source, purpose
Uncertainty	Range or data quality class per figure

One completed record, for the example's option C, shows the profile working; its provenance classifications are illustrative, like its values. Outcome: one thousand documents per batch from the production traffic cohort; success is extraction accuracy at or above 98 percent, audited on a 5 percent sample; all compute, energy, and financial cost incurred by retries and failed cases remains in the reported numerators, while Y counts only cases meeting the success criterion; abstentions route to human review and sit outside Y, with that boundary stated; latency under two seconds at 99 percent availability; contract v1, effective July 2026. Case flow per batch: 1,100 eligible cases presented, 1,080 attempted automatically, 1,000 succeeded, 55 abstained to human review, 25 failed after retry, 20 routed to humans without attempt; successful automation coverage 1,000 of 1,100, roughly 91 percent (attempt rate 98.2 percent; conditional success rate 92.6 percent), with the human review boundary costed separately and disclosed. System boundary: automation subsystem only. Dual denominator at that coverage: \$9.55, 7.45 kWh, and 1.12 kg per thousand eligible cases, before separately accounted human review costs, reported beside the per success figures above. Financial: \$10.50 per thousand, decomposed across compute, energy, and amortized share; no policy or claim expenditure. Energy: 8.2 kWh per thousand, metered. Capacity: 45 kW provisioned against a daytime window of one hour, facility headroom disclosed. Boundary: the named cleaner region, hourly resolution. Emissions method: hourly attributional under the stated factor set, 1.23 kg per thousand, modeled from metered energy. Embodied: excluded, and said so. Instruments: none held. Derived results: none computed. Decision purpose: internal transition planning. Carbon value: internal scenario range for transition planning, \$0 to \$250 per tonne, US dollars, July 2026, selected by management for sensitivity testing against the hourly attributional basis; external market and policy values are treated as contextual inputs, not as interchangeable estimates. Uncertainty: emissions factor banded at plus or minus fifteen percent. Alignment of Outcome Contracts and full population of required fields is comparability's necessary condition, not its sufficient one; a standards process would still supply the normative machinery, canonical field names and units, controlled enumerations, missing value semantics, aggregation rules, and conformance tests. The profile is the specification from which that interoperability standard can be developed, and it is offered as such. The case flow row is the profile's disclosure against gaming: a system can flatter every per outcome intensity by abstaining on difficult cases, so successful automation coverage travels with every per outcome figure, and wherever it is below 100 percent, impact per eligible case is a required companion figure, rather than merely derivable, so that automating only the easiest cases cannot masquerade as efficiency. Where human escalation is material, two records serve: an automation profile with escalations outside the boundary, and an end-to-end service profile counting all completed cases in Y with review costs and, where measurable, review resources inside.

8. What the answer changes

For measurement, the change is the stack: report compute through a vector of financial, energy, and emissions intensities, each normalized to a defined, quality-adjusted outcome, each with its boundary disclosed. For management, the change is the decision layer: adopt one of Section 5's architectures, or a disclosed organizational shadow price, and say which rule is in force, so that the judgment the ledgers make visible has an owner. For capital allocation: when a carbon value enters a financial model, a hurdle adjustment, a scenario stress, a provision, it is a borrowed instrument, and its source, purpose, boundary, and sensitivity belong in the disclosure beside it, because an undisclosed λ applied to an unnamed basis b for an unstated purpose inside a capital model is the governance-as-arithmetic failure of Section 3 in miniature.

The scope, stated as plainly as the claims. This paper does not propose a universal measure of intelligence; it does not select a universal carbon price; it does not treat attributional and consequential emissions as interchangeable; it does not propose to replace GHG Protocol accounting, whose revision it takes as its case study; it does not supply a complete lifecycle assessment method; its reporting profile is a minimum proposal from which a standard can be developed, not an adopted standard; and its worked values illustrate mechanics, estimating no workload. Critique that lands inside these boundaries advances the work; critique aimed outside them addresses a paper this is not.

The coexistence of financial and physical metrics in the field's daily practice is not evidence that conversion is impossible; it is evidence that no existing conversion has displaced the information the separate measures carry. And the participants already behave purpose-specifically, whatever the metrics debate concludes: the largest corporate buyers' holdings rate high risk on the published assessment (Trencher et al. 2024); roughly two thousand firms set internal prices frequently above every observable market (CDP 2023); removals transact at multiples of the reduction price (Ecosystem Marketplace 2025); and the field itself runs two metrics side by side while asking which one is right. The practice has outrun the theory. This paper's contribution is the theory the practice implies.

The diagnosis carries its own boundary conditions, dated on purpose. A universal, enforceable, sufficiently comprehensive carbon price, one number, one authority, most emissions covered, would supply a canonical conversion for compliance cost and materially narrow the problem this paper describes. It would not retire it: the basis behind any figure would still need disclosing, compliance valuation would still differ from damage estimation, abatement ordering, and internal planning, and the energy and capacity ledger would remain untouched, because power constrains a clean grid too. Full scalarization would require more than a price: a shared authoritative emissions basis, a shared decision objective, and a universally applicable value for that objective, a condition nowhere met for carbon today. Nothing in mid-2026 resembles even the narrower condition; the paper is written for the world in which neither holds, and says so.

The surest evidence that no canonical carbon price governs AI compute is that rational operators can face the same tonne, apply different defensible prices to it, and make different decisions without making an arithmetic error. The discipline already carries the necessary measures in practice; what it lacks is the standard that makes their coexistence intelligible.

Notes on sources

Factual claims are stated as of 29 July 2026 and verified against the project's research base on that date. The founding question is closely paraphrased from the cited keynote recording; because the paper paraphrases rather than quotes, the citation is to the recording as a whole, a deliberate choice rather than an omission. Cross-currency price levels are reported in their native currencies and are not numerically ranked against dollar thresholds anywhere in the text. The worked example's values are constructed solely to demonstrate mechanics and estimate nothing. Time-sensitive figures (market prices, consultation status, capital-spending estimates) carry their dates in text; readers after 2026 should treat them as dated by design. Version history: v0.1 drafted July 2026 to a governing outline from a verified research base; v0.2 through v0.12 absorbed ten rounds of structured editorial review and revision, conducted with large language model assistants, with dispositions logged item by item in the companion revision record.

References

Akerlof, G. A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84(3): 488-500.

Haya, B. K., et al. Voluntary Registry Offsets Database, v2026. Berkeley Carbon Trading Project, University of California, Berkeley. Accessed 28 July 2026.

CDP. 2023. *Putting a Price on Carbon: The State of Internal Carbon Pricing by Corporates Globally.* CDP Worldwide.

Congressional Research Service. 2025. *Social Cost of Greenhouse Gases: Background and Recent Developments* (IF12916).

Ecosystem Marketplace. 2025. *State of the Voluntary Carbon Market 2025.* Forest Trends, 29 May 2025.

Electricity Maps. Grid carbon-intensity data by region and hour. app.electricitymaps.com. Accessed 28 July 2026.

EPRI (Electric Power Research Institute). 2024. *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption.* May 2024.

E&E News. 2026. Reporting on EPRI's updated data-center electricity projections (9 to 17 percent of US electricity by 2030), February 2026.

European Commission. 2026a. CBAM quarterly reference prices, Q1 and Q2 2026 (Q2 reference 75.28 per tonne). ec.europa.eu.

European Commission. 2026b. Proposal for reform of the EU Emissions Trading System (Phase 5 framework), presented 17 July 2026. ec.europa.eu.

ICE (Intercontinental Exchange). 2024-2026. EUA futures settlement levels (2024 annual average near 65; July 2026 near 80, December 2026 contract), as compiled in trade reporting by IndexBox and GMK Center, July 2026. Accessed 28 July 2026.

EPA (US Environmental Protection Agency). 2023. *Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances* (EPA-HQ-OAR-2021-0317), December 2023.

FinOps Foundation. 2024. *FOCUS: The FinOps Open Cost and Usage Specification.* focus.finops.org. Accessed 28 July 2026.

FinOps Foundation. 2026. FinOps X 2026 opening keynote, San Diego, June 2026. Recording, youtube.com/live/W4g8xvduIbs. Accessed 28 July 2026.

GHG Protocol. 2025. *Scope 2 Standard Revision and Electricity Sector Consequential Accounting: Public Consultation Materials* (consultation opened October 2025). ghgprotocol.org.

GHG Protocol. 2026. Scope 2 consultation close and process updates (consultation closed 31 January 2026; over 1,000 submissions from 39 countries; revised draft and second consultation announced for 2026). ghgprotocol.org. Accessed 28 July 2026.

Goldman Sachs Research. 2026a. "Tracking Trillions: The Assumptions Shaping the Scale of the AI Build-Out." Goldman Sachs Global Institute and Global Investment Research. goldmansachs.com. Accessed 28 July 2026.

Goldman Sachs Research. 2026b. 2026 hyperscaler capital-expenditure outlook (roughly $527 billion) and US data-center power-demand analyses, January-May 2026. goldmansachs.com. Accessed 28 July 2026.

Hanafy, W. A., et al. 2023. "CarbonScaler: Leveraging Cloud Workload Elasticity for Optimizing Carbon Efficiency." *Proceedings of the ACM on Measurement and Analysis of Computing*

Systems.

Harvard Environmental and Energy Law Program. 2026. The Social Cost of Greenhouse Gases: regulatory tracker, current to January 2026. eelp.law.harvard.edu. Accessed 28 July 2026.

Lempert, R. J., S. W. Popper, and S. C. Bankes. 2003. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis.* RAND.

Linux Foundation. 2026. "Linux Foundation Announces the Intent to Launch the Tokenomics Foundation to Establish Open Standards for AI Cost Management." Corrected release, 3 June 2026.

Litowitz, A., N. Polson, and V. Sokolov. 2026. "Photons = Tokens: The Physics of AI and the Economics of Knowledge." arXiv:2603.06630.

Liu, Q., Z. Chen, and S. X. Xiao. 2022. "A Theory of Carbon Currency." *Fundamental Research* 2(3): 375-383.

McKinsey & Company. 2025. "The Cost of Compute: A $7 Trillion Race to Scale Data Centers." April 2025 (cumulative AI-driven data center investment estimated near $6.7 trillion through 2030).

Neff, D. 2026. "Big-T Notation." Published by the Tokenomics Foundation as an open, vendor-neutral framework, 24 July 2026. tokeneconomics.com/projects/big-t-notation. Accessed 29 July 2026.

NERC (North American Electric Reliability Corporation). 2025. *2025 Long-Term Reliability Assessment.* December 2025.

Pezzey, J. C. V. 2019. "Why the Social Cost of Carbon Will Always Be Disputed." *WIREs Climate Change* 10(1).

Probst, B. S., et al. 2024. "Systematic Assessment of the Achieved Emission Reductions of Carbon Crediting Projects." *Nature Communications* 15.

Shehabi, A., et al. 2024. *2024 United States Data Center Energy Usage Report.* Lawrence Berkeley National Laboratory.

Stechemesser, K., and E. Guenther. 2012. "Carbon Accounting: A Systematic Literature Review." *Journal of Cleaner Production* 36: 17-38.

Swinfield, T., et al. 2026. "Learning Lessons from Over-Crediting to Ensure Additionality in Forest Carbon Credits." *Nature Communications* 17. doi:10.1038/s41467-026-71552-3.

Tokenomics Foundation. 2026. Founding materials and discipline definition. tokeneconomics.com. Accessed 28 July 2026.

Tol, R. S. J. 2023. "Social Cost of Carbon Estimates Have Increased Over Time." *Nature Climate Change* 13: 532-536.

Trencher, G., et al. 2024. "Demystifying the Use of Carbon Offsets by the World's Largest Corporate Buyers." *Nature Communications* 15.

WattTime. 2025. Published analyses of certificate-based accounting effectiveness and consultation commentary on the GHG Protocol Scope 2 revision (finding little if any real-world emissions effect from annually matched unbundled certificates). watttime.org. Accessed 28 July 2026.

West, T. A. P., et al. 2023. "Action Needed to Make Carbon Offsets from Forest Conservation Work for Climate Change Mitigation." *Science* 381.

Wiesner, P., et al. 2021. "Let's Wait Awhile: How Temporal Workload Shifting Can Reduce Carbon Emissions in the Cloud." *Proceedings of Middleware 2021*.

World Bank. 2025. *State and Trends of Carbon Pricing 2025*.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used large language model assistants, including Anthropic's Claude and OpenAI's ChatGPT, for drafting support, source-verification support, critical review, and editorial iteration under the author's direction. After using these tools, the author reviewed and edited all content, independently verified the sources and the claims they support, and takes full responsibility for the content of this publication.

Comments are invited; this is an explicitly provisional working draft.