



BLAVATNIK SCHOOL OF GOVERNMENT

WORKING PAPER

PROVIDING ACCESS TO THE LATEST POLICY-RELEVANT RESEARCH

A REFERENCE ARCHITECTURE FOR IMPLEMENTING E-LEDGERS CARBON ACCOUNTING ON THE FININTERNET

Karthik Ramanna, University of Oxford Blavatnik School of Government

Abhishek Sankritik, Networks for Humanity/ Finternet Labs

Lauren Holloway, E-ledgers Institute

Pramod Varma, Networks for Humanity/ Finternet Labs

Vineet Nair, Networks for Humanity/ Finternet Labs

BSG-WP-2026-64

May 2026

A Reference Architecture for Implementing E-ledgers Carbon Accounting on the Finternet

May 29, 2026

Karthik Ramanna, *University of Oxford Blavatnik School of Government*

Abhishek Sankritik, *Networks for Humanity/ Finternet Labs*

Lauren Holloway, *E-ledgers Institute*

Pramod Varma, *Networks for Humanity/ Finternet Labs*

Vineet Nair, *Networks for Humanity/ Finternet Labs*

Suggested citation: Ramanna, K., Sankritik, A., Holloway, L., Varma, P., & Nair, V. (2026). A reference architecture for implementing E-ledgers carbon accounting on the Finternet. Working Paper.¹

¹ We are grateful for the contributions of Tarebi John, Siddharth Shetty, and workshop participants.

Abstract

We propose a reference technology architecture for implementing E-ledgers carbon accounting on Finternet-compatible infrastructure. We first briefly describe the E-ledgers algorithm and the Finternet for our intended multi-disciplinary audience. We then identify, for software design purposes, the foundational invariants of the E-ledgers system. These invariants – such as mass/energy conservation and full accounting – identify what must be fixed by the architecture and what must therefore be left to entities, auditors, and regulators to determine. Next, we describe the core layers and data objects in the architecture, together with a discussion of data movement, visibility, and control through a comprehensive example. Finally, we discuss the governance and institutional requirements for such an architecture to be deployed across multiple jurisdictions and enterprise systems.

Keywords: carbon accounting, distributed ledgers, E-ledgers, Finternet, offsets.

1. Introduction

The E-ledgers method applies the logic of financial accounting to greenhouse gas (GHG) emissions and removals. Just as financial accounting tracks monetary costs through production and sale, E-ledgers tracks (net) emissions as products and services move through the economy. The aim is to produce emissions accounts that are accurate, comparable, verifiable, and suitable for arm's-length decision making and contracting.

Under the E-ledgers approach, an entity's direct emissions are calculated and verified at source and then allocated to that entity's products along with the emissions embedded in other inputs, using causal logic akin to cost accounting. When a given product is legally transferred to the next entity in a value chain, the emissions travel from the seller's E-ledger to the buyer's E-ledger as an indelible informational property of that product, akin to inventory accounting. Removals are broadly treated in a similar

manner, except that there is no causal logic to allocating removals to products, and the allocation is a purely commercial decision. The E-ledgers system enables all material emissions and removals in value chains to be calculated and assured once and only once. The transfer of emissions and removals through value chains as distinct informational elements of inventory necessitates a technology architecture that allows for relevant data to be trusted by unconnected entities separated by many layers in value chains while preserving the legitimate proprietary information of all those entities.

This paper specifies the minimum reference technology architecture required to operationalize E-ledgers across value chains while preserving the core accounting principles on which the method depends. These principles include counted-once integrity, additivity, auditability, correctability, and the ability to transfer emissions together with products through supply chains. The architecture supplied here is intended to support consistent E-ledgers accounting across different enterprise systems and governance regimes; the architecture is not itself a replacement for those systems and regimes.

The design of the software matters because technical choices can undermine accounting integrity. For instance, a system may sacrifice reliable accounting outcomes under the guise of computing efficiency. Further, problems arise if emissions data becomes separated from the products it belongs to, if historical records can be overwritten instead of transparently corrected, or if reporting is disconnected from the underlying transaction history. Architectural design is therefore not a secondary technical issue, rather, it is necessary for institutional credibility of E-ledgers.

The architecture problem is inherently interdisciplinary. The architecture must work for software systems, accountants, auditors, regulators, and of course businesses using different enterprise services and operating in different jurisdictions. At the same time, architecture should not hardcode legal or policy decisions that properly belong to regulators, standard-setters, auditors, and market participants. This paper therefore

distinguishes between immutable accounting principles that the architecture must preserve and institutional choices that remain open to different jurisdictions, governance frameworks, and ERP products.

The paper specifies the architectural components needed to support product-linked lifecycle events, append-only corrections, aggregate reporting, differentiated evidence access, enterprise integration, and support for both E-liabilities and E-assets. The E-liability side of the framework is developed in greater operational detail because product-level emissions accounting through production, allocation, and legal economic transfer is currently the more mature part of the model. E-assets remain within the architecture as a distinct but linkable accounting object class.

This paper is complementary to the E-ledgers proto-standards and related working papers. Those materials establish the accounting logic, institutional rationale, and broader implementation context. This paper focuses specifically on the technology architecture required to make the model operational.

2. Background

2.1 The E-ledgers method

The E-ledgers method applies the logic of financial accounting to emissions and removals. Just as cost accounting tracks monetary cost through inventory, production, and sale, E-ledgers tracks emissions as products move through supply chains (Kaplan and Ramanna, 2021; Kaplan and Ramanna, 2025). The goal is to produce emissions information that is as accurate, timely, comparable, and auditable as financial information, and to do so in as close to real time as practicable.

Under the method, the entity that directly emits GHGs records those emissions once at source. When products are sold, the emissions embedded in those products transfer with the sale. The receiving entity then adds its own locally incurred emissions,

allocates the combined total to its own outputs, and transfers those onward at the next sale.

The same emissions are therefore not recalculated independently at every stage of the supply chain. Instead, they are measured once at source, carried forward, and reallocated as products are transformed and sold. In this respect, the method operates similarly to a value-added system.

There are two sides to an E-ledger: E-liabilities and E-assets. E-liabilities represent emissions associated with products, activities, events, or conditions. E-assets represent qualifying removals that are indefinitely sequestered and under the exclusive legal use of an entity (Kaplan, Ramanna, and Roston, 2023).

2.2 Finternet as a coordination infrastructure

Product-level emissions accounting only becomes useful at scale if emissions records can move reliably across firms, supply-chain tiers, and jurisdictions. Companies use different enterprise systems and operate under different legal and reporting environments. A workable architecture therefore needs a way for emissions records to move between organizations without requiring all participants to adopt the same software platform or centralized infrastructure.

Finternet is a network-of-networks model for interoperable digital assets, credentials, and transaction infrastructure intended to support this kind of coordination across institutions and jurisdictions (Carstens and Nilekani, 2024). Within the context of E-ledgers, the accounting method defines what should be tracked, while interoperable infrastructure defines how those records can be represented, transferred, verified, and governed across organizations.

Several features of Finternet align directly with these requirements. Product-linked digital records (“product tokens”) can carry emissions information together with commercial transaction data. Control rules set by a universal governance layer can

determine how emissions records are transferred, updated, or corrected while preserving counted-once integrity. Digitally verified credentials issued by intermediaries like registered auditors allow emissions claims to be trusted as they move downstream, without requiring every participant to disclose all underlying evidence. Interoperable ledger infrastructure can support coordination across enterprise boundaries without forcing all participants onto a single centralized system.

These features are not separate from the accounting model itself. Product-level emissions accounting requires a way to link emissions to products, transfer those records across organizations, preserve audit history, support corrections, and maintain trust between participants operating in different systems and jurisdictions. Finternet-compatible infrastructure provides one possible foundation for meeting those requirements.

2.3 Scope of the implementation architecture

The paper specifies the minimum technology architecture required to operationalize E-ledgers across interoperable enterprise environments. It focuses on preserving the accounting rules the method requires rather than prescribing a single technology stack, governance framework, or deployment model.

Some requirements are treated as fixed (“invariants” set out in §3). These included counted-once recording, append-only correction logic, mass/energy conservation across transfers and allocations, and the maintenance of auditable event history. Other questions are deliberately left open because they depend on legal, regulatory, institutional, or market decisions. These include ERP integration approaches, the use of public versus private ledger infrastructure, acceptable fallback data sources, evidentiary standards for a legal economic transfer, disclosure obligations, and verifier eligibility.

The result is a reference architecture intended to support interoperability, auditability, and institutional portability, while remaining adaptable across different regulatory and technological environments.

3. Invariants

To implement the E-ledgers method, a small core of prescriptive requirements must be enforced regardless of the technology system, jurisdiction, sector, and participants.

These are the invariants of the proposed architecture. The invariants ensure that any system deployed is internally consistent, externally credible, and interoperable.

Functions like an auditor attesting emissions assertions, a counterparty accepting a transferred position, a regulator requiring aggregated reports, and a downstream party being able to re-use upstream data without reopening it are built atop these invariants.

This section identifies seven invariants. (Table 1 provides a summary.)

3.1 Counted-once recording

The entity that directly generates GHGs from its operations records those emissions once, at the place they arise. Downstream buyers do not mint fresh direct-emissions records for the same physical emissions, reflecting the fact that the emissions are created only once. They receive a transferred embedded-emissions position when they buy the product to which those emissions have already been allocated.

3.2 Transfer-linked movement

Embedded emissions move with the product and with the legally relevant transfer of that product. The architecture links the emissions position to the product token, which is the digital counterpart of the commercial object being bought, rather than to a detached certificate or a free-floating data item. The buyer receives the emissions position in the same transaction by which it receives the underlying product.

3.3 Full allocation under causal logic

All recognized emissions arising from economic activity must ultimately be allocated to outputs following a causal logic. The architecture must preserve the relationship between the physical and economic cause of the emissions and the output to which they are assigned. The allocation method must be explicit and reviewable for an auditor to test. The architecture therefore carries not only the resulting product-level emissions number but also a reference to the methods under which that number was produced.

The E-liability proto-standard (Ramanna et al., 2025) identifies a few (short-run) situations in which full allocation to products or to capitalized equipment is infeasible, generating a temporary unallocated-emissions account: any quantity of recognized input emissions or asserted direct emissions that has not been bound to intermediates or to an output product token at the close of an allocation period is automatically captured in this account. A non-zero balance surfaces immediately to the auditor as a flag of incomplete allocation. The mass/energy conservation equation, which states that the sum of input-side recognized emissions equals the sum of output-side allocated emissions, is therefore always machine-readable, and is always either in balance or visibly flagged.

3.4 Conservation

Once recorded, emissions cannot be lost, duplicated, or reassigned without a trace, consistent with the first law of thermodynamics. As a direct consequence, various controls are required. The product token must be quantity aware, i.e., it cannot be invariant to the mass of emissions it represents. The architecture must also preserve dependency links between outputs and the recognized quantities of inputs used to produce them. The allocation step must further produce a balance, in which total assigned emissions across outputs, capital equipment, and retained inventory must equal the sum of transferred-in and locally incurred emissions. Ensuring mass/energy conservation is always maintained also suggests the need for gross-history retention.

For instance, where an entity nets an E-liability against a recognized E-asset, retaining the gross original position allows for robust bookkeeping in cases where the asset-side position is later impaired.

3.5 Append-only correction

A later correction does not overwrite the original record. The system adds a new record stating that the earlier record has been superseded but the earlier record remains visible. This is the difference between visibly correcting a ledger and secretly editing an earlier record as though the edit never occurred. Product-level emissions accounting is temporal. A company may sell in March on one attested figure and receive a corrected attestation in June. The architecture must serve both the auditor who needs the March number and the manager who needs the June number.

3.6 Layered disclosure

Different participants need different views of the same data. The architecture separates what must be shared for accounting integrity from what should remain under controlled access for competitive purposes. A counterparty receiving a product needs only the relevant verified emissions position transferred with the product to accept and record it accurately. An auditor requires deeper visibility into source measurements, methodology, and allocation logic to verify accuracy and compliance. A regulator may have statutory authority to access all details, including transaction-linked and aggregate totals. The public may see only consolidated entity-level statements. Each access class receives only what its function requires, as prescribed by law and determined by the data owner. This is a structural requirement on the architecture itself, not a feature added downstream. For instance, auditability cannot be made to depend on public disclosure of sensitive enterprise data.

3.7 Interoperability

Interoperability is a requirement for different systems and different jurisdictions to understand what a number refers to, what method produced it, what assurance attaches to it, and how it can be reused without recalculating the whole chain. That requires a schema or a common structure for product identifier, product quantity, emissions value, allocation method (i.e., how product-level emissions were calculated), and verification status. Interoperability also requires an exchange logic, a transport model for the cross-boundary movement of structured records.

4. Reference technology architecture

In this section we cover how the invariants discussed earlier are implemented in the reference technology architecture. This section sets out the layers on which the records sit, the core record classes themselves, and the rules that govern movement, visibility, and control between participants. Functional roles and authority rules are detailed in Appendix A, and the principal terms also appear in the Glossary.

4.1 Three layers of record

There are three layers of record in the architecture: the E-ledger itself; the shared layer; and the E-ledger *statement*. (Table 2 provides a summary.)

The E-ledger is the entity-level book of account for emissions. It is the environmental analogue of the financial accounting ledger that records the financial impact of events in the enterprise. The E-ledger records the emissions impact of day-to-day operational events, as declared by the entity. These events can include purchase orders, goods receipts, production completions, inventory movements, and meter readings. The E-ledger can be constructed in part from existing ERP systems that already capture underlying business events.

The shared layer is the cross-organizational coordination layer that records product-linked state, transfer events, attestation references, and other elements that must be portable and tamper-evident across organizations (these terms are defined in the Glossary). The “shared layer” term is used as the category label, without commitment to any particular technology. It implies the property of being shared and tamper-evident across enterprise boundaries. A distributed ledger is one such technology choice. Permissioned registries and federated databases are others. The architecture is neutral on which substrate is used, provided the shared-layer property holds. The sole requirements from the shared layer are that it is accessible to relevant participants and that information posted to it is tamper evident.

The E-ledger statement is the reporting view derived from the E-ledger’s event history. It is the environmental analogue of a financial balance sheet: with an opening balance, additions and dispositions during the period, and a closing balance. The statement is derived from the underlying E-ledger.

4.2 Core objects

There are three core objects or “records” in the technology architecture proposed. These core objects further reference underlying attestations (made by independent auditors) that are central to the records’ integrity in the architecture. The core objects and attestation linkages are introduced below.

A **product token** is the digital counterpart of a defined commercial item or quantity, such as a batch, lot, shipment line, or unit. It is the central commercial object the architecture transfers across enterprise boundaries. The product token carries essential information about the product (output) that it represents, and it is the carrier through which the linked embedded emissions transfer. Entities create output product tokens and in turn accept input product tokens through legal economic transfer. Entities may also create *intermediate* output product tokens, i.e., tokens that represent items at intermediate stages of production, which can be then transformed along with further

input product tokens and direct-emissions processes into final output product tokens. These intermediate output product tokens are functionally identical in the architecture to final output product tokens.

Product tokens must be divisible, allowing a fractional transfer of stock to carry the corresponding fractional embedded emissions in mathematically exact form. The product token, as defined in this section, is a unique carrier for E-liabilities and is distinct from an E-asset token, which represents a recognized removal and sequestration position (described in §5).

A **direct-emissions declaration** is a versioned, signed claim about the direct emissions generated by an entity from a production event. A direct-emissions declaration is eventually linked to one or more output product tokens via allocation records (defined below). Direct-emissions declarations move through two basic states: provisional (originated by the producing entity, awaiting attestation), attested (signed by an accredited verifier). Further, a provisional direct-emissions declaration can be “superseded” (replaced by a later unattested version). An unattested direct-emissions declaration is generally internal to the originating entity. Attested declarations cross the shared layer with the verification credential that travels alongside them.

An **allocation record** is the record that binds a defined quantity of emissions, drawn from one or more source emission records, to one or more (intermediate or final) output product tokens under a recognized allocation method. The sources can be the entity’s own direct-emissions declarations, input product tokens transferred in from a tier-1 supplier, and the entity’s own intermediate product tokens. Multiple allocation records may attach to a single output product token, for example where various overhead emissions represented by different direct-emissions declarations are spread across several product lines. Every allocation record must be faithful to the mass/energy conservation invariant introduced in §3.4, so the conservation equation holds across all inputs and outputs to that allocation. Each allocation record references a supporting record called the **allocation method attestation**: a signed record from the auditor

stating which methods were used, for what quantities, drawn from which sources, and bound to which output product tokens.

Allocation records create **dependency links** between output product tokens and the relevant emissions sources. A dependency link is queryable in both directions. Read from an output to its inputs, the link supports uncovering the provenance of an output product token's embedded emissions. Read from an input to the outputs that drew from it, the link is an additional investigation primitive. An auditor or a regulator reviewing an upstream entity's records, or a counterparty raising a dispute, can enumerate the population of output product tokens that received transfers from a specific source (e.g., a fraudulent direct emissions declaration).

4.3 Movement, visibility, and control

The architecture above introduces a deliberate design feature that merits further discussion: Product tokens can “move” only one tier in a supply chain. When legal economic transfer of a product occurs, and a product token transfers from seller to buyer, the buyer does not transfer that token unaltered to its own customers. Rather, the buyer links its input product tokens to its own new output product tokens. The new tokens reference the input token's emissions through dependency links, preserving mass/energy conservation. Each subsequent tier behaves identically.

One-step movement through these links is a deliberate design choice grounded in protecting privacy. In a competitive market, a seller's process detail and a buyer's downstream use of that seller should be visible only to parties whose functions require that visibility. One-step movement ensures that each tier sees only what it directly needs. The alternative, where a token traverses many tiers under its original identifier, could be computationally intensive and could result in every tier having to expose its proprietary details to all further downstream tiers in the chain. The technology architecture here doesn't rule out that alternative; it simply picks a computationally feasible solution that meets the necessary criteria.

Of course, an entity always owns the choice of how much visibility to allow on top of one-step movement. Several arrangements are common in practice. A vertically integrated supply chain with a single legal owner of several production stages may permit full traversal across many supply-chain tiers, with the product token passing through each stage and emissions positions visible at every step. An arm's-length supply chain typically operates at the stricter end: each tier sees only the immediately upstream attested figure, with internal production detail kept for the independent auditor's eyes only. A sectoral consortium may sit between the two, with aggregated visibility across consortium members and stricter separation from non-members. Each pattern is a configuration choice on top of the same base architecture. The choice itself should sit with the relevant entities, to be determined by market forces and regulation.

For entities that elect to keep visibility tight, three operational properties of the architecture are relevant. First, different participants get different views of the same record through differentiated access classes for counterparties, auditors, regulators, and public viewers. Second, an upstream attestation can be relied on by the recipient without disclosing the underlying evidence to that recipient: the credential is signed, and the evidence stays at source, indelibly. Third, where deeper access is later legitimately required, in a dispute or regulatory inquiry, the system points reliably to the underlying records. The detail of access classes, deployment models, and cryptographic tooling is in Appendix B.

5. How the architecture operates: a stylized case study

This section traces the architecture through the example of a three-tiered supply chain. In our example, the first company in the supply chain, SteelCo, produces 100 units of coil carrying 200 kg of direct emissions. The second company, FabricatorCo, buys 40 units of that coil, fabricates 20 output components using its own production emissions, and sells 10 of those components to the third company, GreenCo. GreenCo holds an E-

asset that it applies against the embedded position of the acquired components, as part of its premium low-emissions offerings, and then resells the product to the end consumer. In what follows, we describe how each of these steps unfolds in the architecture introduced earlier.

In the first step, SteelCo, produces 100 units of coil carrying 200 kg of direct emissions. For simplicity, we assume that there are no further input emissions to SteelCo. This step is then associated with SteelCo minting an **output product token** representing the 100 units. Separately, as SteelCo incurs the 200 kg of direct CO₂ emissions, it issues a **direct emissions declaration** to that end. For this record to be trusted by its customers, SteelCo has this declaration record **attested** by its auditor. Next, SteelCo creates an **allocation record** when it binds the direct emissions declaration to the output product token. This allocation record is also verified by an auditor, whose assurance is attached to the record as an **allocation method attestation**. For simplicity, we assume a mass-based allocation method, so each unit of coil is assigned 2 kg of CO₂.

The architecture is agnostic on storage of the internal records, as long as they are accessible by relevant parties and tamper evident. For instance, the direct emissions declaration and the allocation record, together with their respective signed attestations, may sit in a private database, in a credential wallet, or on a public ledger as a zero-knowledge proof credential. What the architecture requires is that these records be retrievable by version, with indelible notes on any editing or attestation history, upon legitimate request (e.g., by an auditor or regulator, or upon court order).

SteelCo's output product token must be capable of being linked to its customers. The architecture also specifies that the product token be divisible (up to eighteen decimal places) to accommodate for potential divisibility of the production units to which that token attaches. So, when SteelCo sells 40 of its 100 coil units to FabricatorCo, the output product token divides cleanly: FabricatorCo receives a sub-token covering exactly 40 units (and 80 kg of CO₂), while SteelCo retains 60 units (and 120 kg of CO₂).

In the architecture, the product token division and transfer are authorized by the “legal economic transfer” of the underlying product (here, the coil units), as evidenced by a fiscal record, such as an invoice, recognized in the relevant jurisdiction (Ramanna et al., 2025). That fiscal record becomes the **transfer-evidence reference** (more details in Appendix C) on which the token division and transfer, or change in claim, are predicated. The transfer-evidence reference can be subject to verification and can embed an attestation similar to those embedded in direct emissions declarations and allocation records. Fiscal records vary across jurisdictions: the transfer-evidence reference can be programmed to accept whichever fiscal records the relevant jurisdiction recognizes as admissible.

In the second step, FabricatorCo, having acquired 40 units of the coil carrying 80 kg of CO₂, produces 20 finished components for its customers. This production necessitates FabricatorCo to mint an output product token of its own. Further, FabricatorCo’s production incurs direct emissions of 20 kg of CO₂, so it records a direct emissions declaration to that amount which it has attested by its auditor. Next, it creates an allocation record, also subsequently attested, to allocate its direct emissions and the input product (sub)token, i.e., a fraction of the original product token, it received from SteelCo to its output product token.

Since FabricatorCo has 80 kg of input emissions and 20 kg of direct emissions, it must allocate the total of 100 kg to its 20 finished components (assuming no other activities and outputs). Assuming uniform allocation, each finished component is apportioned 5 kg of CO₂, which is represented in FabricatorCo’s output product token.

The architecture ensures that the assurance status of the upstream emissions, embedded via reference in the (sub)token transferred from SteelCo to FabricatorCo, carries to the recipient. This allows the recipient, FabricatorCo, to rely on the emissions value in the transferred (sub)token without having to reverify it. The bidirectional **dependency link** in SteelCo’s output product token permits FabricatorCo’s auditors (or other legitimate users such as regulators) to trace back emissions to ensure legitimate provenance: for

instance, supplier compliance with invariants such as causally based full-allocation and mass/energy conservation. Likewise, SteelCo's auditors (and other legitimate users) can query where SteelCo's emissions have been apportioned downstream. These queries are read-only: they cannot be used to alter emissions values once allocated and transferred.

In the third and final step of our stylized example, FabricatorCo sells 10 of its 20 output components to the third company in the value chain, GreenCo. Earlier, FabricatorCo had created an output product token representing its 20 components, with a total emissions value of 100 kg of CO₂ and a per-component value of 5 kg. So, the sale of 10 components to GreenCo is accompanied by a transfer of a sub-token representing 50 kg of CO₂ (at 5 kg per component), with the balance 50 kg of CO₂ being retained by FabricatorCo.

The core objects in the architecture (direct emissions declarations, input product tokens, allocation records, and output product tokens) are manifestations of activities on an entity's E-ledger. Thus, data from the core objects during an event period can be consolidated into an **E-ledger statement**, providing an entity-level reporting view derived from product- and transaction-level records. For instance, FabricatorCo's E-ledger statement, based on the information described thus far, can be prepared from the core objects as shown in Table 3.

In our stylized example, the final company in the value chain, GreenCo, holds an E-asset worth 30 kg of CO₂ that it applies against six of the 10 acquired components, selling these as "net zero" products to premium-paying end consumers. The remaining four components are sold as lower-end products at their full emissions value of 5 kg of CO₂ per component.

The introduction of the E-asset in this part of the example allows us to discuss the **E-asset token**, previously mentioned but not described in detail. An E-asset token represents an audited quantity of E-assets, which refer to distinct and transactable amounts of indefinite-duration sequestrations of already generated CO₂. Details of E-

assets are described in the E-asset proto-standard (Ramanna et al., 2026). An E-asset token represents a unique quantity of E-assets available for netting against E-liabilities, so use of such a token prevents duplicative use of those E-assets elsewhere.

An E-asset token sits on the asset side of an entity's E-ledger as a recognized position waiting to be netted or sold. When an entity *nets* its E-asset token against the E-liabilities in one or more of its output product tokens, the E-asset token moves from the **held state** to the **applied state**, which is simply a flag to ensure the E-asset token is not reused. Netting an output product token with an E-asset token, allows the output product token to transfer a lower net E-liability to that entity's customer when the underlying product is sold. (The output product token still maintains the gross E-liability position for recordkeeping purposes.) If the E-assets represented by an already netted E-asset token become impaired, the architecture requires that the netting entity records a direct emissions declaration to the extent of that impairment. These emissions, implied from the asset impairment, must then be allocated to the entity's future output product tokens, preserving mass/energy conservation. Note that downstream customers who have already purchased a low-emission product attributable to E-asset netting that subsequently becomes impaired are not affected by that impairment (Kaplan, Ramanna, and Roston, 2023).

Returning to our example, as GreenCo is selling six of its 10 acquired components as premium "net zero" products and four as regular products, it first mints two output product tokens, one for each product class. Then, it creates an allocation record to link the 50 kg emissions in its acquired input product (sub)token from FabricatorCo to its two output product tokens, with 30 kg assigned equally across the six components in the premium category and 20 kg assigned equally across the four components in the ordinary category. Finally, GreenCo nets its 30 kg E-asset token against the six components in the premium category, yielding a net zero E-liability for that output product token and switching E-asset token flag from held to applied.

As GreenCo's products, both premium and ordinary, are sold to end consumers, the corresponding output product tokens are transferred to the consumption E-ledgers of the jurisdictions where those sales take place (see Kaplan and Ramanna, 2025).

Nearly every step in the stylized example above represents events and entails data already captured in modern entities' records systems. For instance, SteelCo's 200 kg of direct emissions can be calculated from the stoichiometry of its chemical processes and corresponding data on the quantities of chemicals deployed in those processes, as inferred from inventory records. Likewise, its sale of 40 coil units to FabricatorCo is already recorded in invoices. The technology architecture developed in this paper maps these existing events and data onto E-ledgers while preserving the E-accounting invariants.

Indeed, we conjecture that much of carbon accounting's failures today can be addressed through improved data integration, as much of the data for robust carbon accounting already sits across uncoordinated systems, spreadsheets, and PDFs. The architecture presented here can work with existing enterprise stacks, drawing and imputing primary data from events already recorded. The architecture is agnostic to heterogeneous enterprise environments; it sustains interoperability through common accounting semantics, recognized transfer structures, standardized assurance, and compatible exchange logic across systems.

6. Governance

The architecture described in this paper is itself a form of governance that holds any software deployment of the E-ledgers system consistent with the invariants. Beyond this governance layer, the functioning of E-ledgers via this or other architectures will, at a minimum, require two additional governance layers: technology-related governance and accounting-related governance. In this section, we briefly describe these complementary governance layers to offer sight of what is needed to make E-ledgers

practicable and to implicitly make the argument that this comprehensive governance ecosystem is feasible.

Technology-level governance covers issues such as who maintains the shared layer for a given jurisdiction or sector, which schema versions for product characteristics are recognized, and which token managers are accredited to run state transitions. While these choices must ultimately be ratified by the jurisdictions where E-ledgers data are being accepted and used for legally enforceable commercial transactions, the choices can be made bottom-up by industry-level bodies or even by transacting counterparties, particularly those with considerable market power in their value chains.

Accounting-related governance covers issues such as materiality and measurement thresholds for direct emissions recognition, rules for recognition of removals as E-assets, acceptable allocation methods (i.e., acceptable interpretations of causal logic), evidentiary rules to document legal economic transfer, assurance standards, eligibility criteria for auditors, and penalties for non-compliance. In financial accounting, these issues are adjudicated by standard setters, capital-market regulators, and statutory law. In carbon accounting, the proto-standards produced by the E-ledgers Institute are first attempts at addressing these issues. Given the potential uses of E-ledgers in revenue-raising public-policies such as carbon border tariffs, we expect such issues to be resolved by public Emissions Accounting Boards and associated statutes in relevant jurisdictions.

In the near-term at least, we advocate that the choices in technology- and accounting-related governance layers remain principles-based rather than highly prescriptive, reflecting that decarbonization technologies and approaches, and the measurement thereof, are themselves fast-evolving. Moreover, as the E-ledgers proto-standards note, onboarding a system like E-ledgers within and across jurisdictions will require a grace period for learning and improvement in value chains, so the strictness of accounting and auditing rules, and associated enforcement, should ideally be ramped up over time.

Note that the reference technology architecture introduced here is itself invariant to the evolution anticipated above. Indeed, technology evolution is built into the architecture so that, for instance, a regulator can recognize a new measurement approach, or a standard setter can publish a refined allocation rule, without requiring every prior token to be re-issued. Each record carries the version characteristics under which it was produced, the shared layer carries records produced under multiple concurrent versions, and receiving parties can easily determine which version applies to a given token. Regulators can specify the degree of backward-compatibility permissible by specifying which version tokens remain fungible, for what applications (products, processes, and jurisdictions), and for how long. The functional benefit of this adaptability is that a jurisdiction, sector, or entity can, for instance, immediately adopt an improvement for one (highly material) product line without radically disrupting long-planned transactions in all other product lines.

7. Conclusion and open questions

Several questions remain open at this stage. Some are external to the architecture and properly belong with policymakers and market participants. Others concern the implementation of this architecture and should be developed by the relevant regulators, counterparties, and intermediaries.

- **Legal questions:** The architecture assumes a working answer to several legal questions without itself settling them, including: what rights auditors and regulators have to deeper evidence layers behind a credential; how these rights travel across borders; the liability of infrastructure operators for publishing and maintaining information that others rely upon; and how tokenized or digitally represented records interact with existing property and contracts law.
- **Policy questions:** The architecture is compatible with, and indeed agnostic to, multiple policy regimes, and several questions about which regime applies

remain external, including: whether public-ledger publication is required; what fallback sources are accepted for absent primary data; what assurance phase-in is appropriate for new sectors or jurisdictions; who may access aggregate versus transaction-linked data; and how asset-side positions may be used or regulated.

- **Enterprise integration:** The architecture maps everyday commercial events onto the E-ledger, but the precise API surface for integration with ERP systems is left to implementation intermediaries. The depth at which integration is presupposed will shape both deployment cost and what auditors are asked to attest. Enterprise integration providers are best placed to lead on these matters, with regulatory approval eventually needed. Hardware devices that can be integrated into this stack is another matter that must be deliberated, with an aim to arrive at standards on reliable automated emission recordings.
- **Practical operation of the layered governance structure:** The architecture presupposes two complementary governance layers related to technology and accounting choices. How those two layers are developed and how they interact with each other and the architecture remains a question outside the scope of this framing document.

This paper translates the E-ledgers method into a deployable reference technology architecture. Because the architecture can build from existing enterprise and legal infrastructure, even rapid scaling, commencing imminently, is feasible. At scale, it changes what is operationally possible in product-level emissions accounting. For instance, accounting and audit costs today compound at every tier of a supply chain as each downstream company and auditor (where even used) re-calculates upstream data; under this architecture, an attested record issued at source travels downstream for seamless, fungible use. Product-level emissions data becomes comparable across producers and across jurisdictions, providing the operational precondition for regulatory regimes such as border carbon tariffs. Moreover, incentives align with measurement, mitigating the need for moral suasion to drive decarbonization: a

producer's investment in cleaner operations flows through to the embedded position of every downstream product that draws on its outputs, rewarding cleaner production. Adoption is self-reinforcing, since each new participant with attested data improves the accuracy of systemwide data and thereby improves any endogenous commercial returns to net emissions reductions.

Glossary

Term	Definition
Accounting ledger	An entity's enterprise system of record for operational and financial events, including, purchase orders, invoices, goods receipts, production completions, and inventory movements. The financial-accounting analogue of the E-ledger; in this paper it appears only in that comparative role and is not the primary architectural label for the entity-level emissions ledger.
Allocation method attestation	A signed record from the auditor stating which allocation method was used and on what scope. Emissions allocations reference this attestation directly so the auditor can verify both the underlying emissions number and the rule that produced it.
Allocation record	The record that binds a defined quantity of emissions, drawn from a source emission record that is either a direct-emissions declaration or an input product token, to one or more output product tokens under a recognized allocation method. Carries quantity, source reference, target product token, method identifier, assurance status, and allocation method attestation reference.

Append-only correction	The principle that later corrections add new records rather than overwriting earlier ones, preserving the full historical chain. Example: an emissions assertion attested at 1,000 kg in March and revised to 1,080 kg in June produces two records, both visible; March remains historical while June becomes the latest-truth view.
Auditor	The role providing authoritative assurance through attestation. On the E-liabilities side, this includes attestation of direct emissions declarations, allocation methods, and transfer evidence. The auditor produces verified credentials that travel with the relevant records.
Counted-once integrity	The principle that a unit of direct emissions is recorded once at the place it occurs and carried forward through allocation and transfer, rather than re-estimated at each downstream stage.
Dependency link	The structure within an allocation record linking an output product token to a defined quantity of emissions from a direct-emissions declaration or from an input product token, under a recognized allocation method. Queryable in both directions, from output to source and from source to outputs that drew from it. Supports auditor and regulator investigation, and supply-chain awareness. Queries are read-only; post-attestation corrections are handled through a new direct-emissions declaration at the responsible entity, allocated against future outputs.
Direct emissions declaration	A versioned, signed claim about the direct emissions generated by an entity from a given activity or period. Becomes authoritative when an auditor attests it.

E-asset	E-assets are distinct and transactable amounts of indefinite-duration sequestrations of already generated CO ₂ . An E-asset token represents a unique quantity of E-assets available for netting against E-liabilities, so use of such a token prevents duplicative use of those E-assets elsewhere.
E-ledger	The entity-level book of account for emissions – the emissions analogue of an accounting ledger that records the financial impact of business transactions, events, and conditions. Distinct from the ERP systems that capture the underlying business events.
E-ledger statement	The periodic reporting view derived from the E-ledger’s event history, analogous to a financial balance sheet. In the reference technology architecture, the E-ledger statement is derived from the underlying E-ledger; not stored as a separate authoritative record.
E-liability	A distinct liability-side environmental accounting object representing a unit of greenhouse-gas emissions (e.g., kg of CO ₂ , CH ₄) associated with an entity’s activities.
Finternet	A network-of-networks model for interoperable digital assets, credentials, and transaction infrastructure. It is an example of the shared coordination layer on which the E-ledgers architecture can operate.
Governance layers	This paper refers to three governance layers: the architecture itself, the technology-governance layer and the accounting-governance layer. Technology governance covers issues such as who maintains the shared layer for a given jurisdiction or sector, which schema versions for product characteristics are recognized,

and which token managers are accredited to run state transitions. Accounting governance covers issues such as materiality and measurement thresholds for direct emissions recognition, rules for recognition of removals as E-assets, acceptable allocation methods (i.e., acceptable interpretations of causal logic), evidentiary rules to document legal economic transfer, assurance standards, eligibility criteria for auditors, and penalties for non-compliance.

Input product token	A product token an entity has acquired from a supplier; it carries E-liabilities and is consumed in production.
Invariant	A non-negotiable rule, derived from E-ledgers principles, that must hold in every implementation. The invariants in §3 are the conditions any technical architecture must preserve to remain consistent with the E-ledgers method.
Layered disclosure	The principle that different participants – counterparty, auditor, regulator, public – may legitimately receive different views of the same data, based on what their function requires.
Legal economic transfer	A legal economic transfer of a good or service is one that is effected consensually between entities to a transaction, usually for monetary or similar consideration and results, where applicable, in a legally recognized transfer of ownership and control of the underlying unit being transacted. The architecture binds product-token transfer to a transfer-evidence reference that evidences the legal economic transfer. E-liabilities move across E-ledgers with the legal economic transfer of the underlying product.

Output product token	A product token an entity has minted to represent a commercial batch, lot, shipment line, or unit of its own output. The token carries the emissions allocations attributed to it and may be transferred to a buyer. Distinct from an E-asset token, product tokens carry E-liabilities only. Product tokens are divisible to eighteen decimal places, so a fractional transfer carries a mathematically exact share of the embedded emissions. Entities may also create intermediate output product tokens, i.e., tokens that represent items at intermediate stages of production. These intermediate output product tokens are functionally identical in the architecture to final output product tokens. Where the role matters, the paper uses input product token (acquired from a supplier) and output product token (minted by the entity) as distinct terms, but they are identical elements in the architecture.
Reference technology architecture	A structured description of the layers of record, the records themselves (i.e., core objects), their roles, their evidence rules, and authority rules for their movement, visibility, and control that any invariant-compliant implementation must preserve. More abstract than a software product specification, it does not prescribe vendor or technology choice.
Shared layer	The coordination layer that records product-linked state, transfer events, attestation references, and other elements that must be portable and tamper-evident across organizations. A distributed ledger is one technology option providing a shared layer.
Substrate	The underlying technology or digital medium used to record and exchange information across a supply chain. An ERP database, a permissioned ledger, or a public blockchain are each substrate

choices. The accounting result must be the same across substrates.

Token manager	The control logic that governs state transitions on a token: who may originate it, how it transfers, what conservation rules apply, how post-attestation corrections are recorded. In this paper, the token manager enforces the invariants set out in §3.
Transfer-evidence reference	The link between a product-linked transfer and the commercial or regulatory proof point that corroborates it: e.g., an invoice, e-invoice, customs declaration, warehouse release, stamp or fee receipt. Appendix C sets out how jurisdiction-specific records are normalized.
Unallocated-emissions account	A standing record at the closure of each allocation period that captures any quantity of recognized input emissions or asserted direct emissions that has not been bound to an output product token. A non-zero value flags incomplete allocation immediately to the auditor.
Verified credential	A digital attestation issued by an auditor or other trust provider that can travel with a product-linked record and be checked downstream without the recipient inspecting the full underlying evidence set.

Appendix A. Functional roles and authority rules

A.1 Roles

A single organization may act as buyer, producer, seller, inventory holder, or E-asset holder at different points and in different transactions. The architecture distinguishes functional roles rather than organizational identity.

The **producer** creates the first output product token for a given quantity of direct emissions and originates the initial direct emissions declaration against that token. The **seller** transfers the product state and the linked embedded emissions on the basis of a recognized transfer-evidence reference. The **buyer** receives the product-linked state and the embedded emissions allocation from the immediate seller. The **auditor** provides authoritative assurance through attestation of direct emissions declarations and of allocations and allocation methods. The **governance authority** recognizes methods, schema versions, verifier eligibility, and credential rules. The **integration operator** maps enterprise events into the shared layer as a technical function distinct from the substantive accounting or assurance authority. The **E-asset holder** holds, applies, and impairs E-assets under the relevant E-asset standard.

A.2 Authority rules

Origination is distinct from attestation. An issuer may originate a product token and an initial direct-emissions declaration. Attestation requires an accredited verifier. Any later change to effective emissions occurs through a new attested direct-emissions declaration allocated against the next batch, or another explicit state transition recognized by the governance authority.

The token manager enforces the authority rules. It permits transitions that conform to the governance configuration in force and rejects those that do not. The substantive content of the rules, including recognized methods, accredited verifiers, and evidentiary

forms for legal economic transfer, is supplied by the governance authority. The technical layer simply enforces; it does not legislate.

Appendix B. Privacy toolkit

This appendix specifies the technical patterns that satisfy the privacy functional requirements introduced in §4.3. Privacy is treated as a layered design decision rather than a single primitive, because the same architecture is asked to support a wide range of access classes, including counterparty, auditor, regulator, and public, and a wide range of deployment models, including permissioned, public, and hybrid. The patterns below are stated as design options, not as architectural mandates: a deployment selects the patterns appropriate to its commercial context, jurisdiction, and risk profile.

B.1 Information layers

On-ledger records carry token identifiers, emissions balances, state transitions, attestation references, and commitments to evidence packages. They are the minimum globally verifiable layer. Off-ledger enterprise vault records carry detailed allocation calculations, raw measurement data, supplier-specific factors, and internal process documentation. They are disclosed only under selective-disclosure rules. Cross-organizational evidence relays carry commitment hashes linking on-ledger references to off-ledger evidence, allowing later verification without prior disclosure.

B.2 Access classes

Four access classes are defined at this stage, although further access classes can easily be added. Counterparty access carries the embedded-emissions number, the allocation method identifier, and the assurance status, which is enough to take a transferred-in position. Auditor access carries source measurement data, allocation workpapers, and the full transfer-evidence chain, which is enough to test underlying numbers and

methodology. Regulator access carries all data mandated under statutory authority. Public access carries entity-level or product-class statements where required by disclosure laws. Each access class receives only what its function requires.

B.3 Evidence and attestation portability

A recipient must be able to rely on an upstream attestation without the issuer disclosing the underlying evidence to that recipient. Verified credentials provide the mechanism. An auditor signs a credential stating the assured number and the method. The credential travels with the product token. Downstream parties verify the auditor's identity and the credential's integrity without seeing the underlying data. Where deeper access is later legitimately required for further audits, dispute, or regulatory inquiry, the system points reliably to the underlying records, even where they remain under controlled access until that point.

B.4 Cryptographic tooling

Commitment schemes publish cryptographic commitments on-ledger to bind evidence without revealing it. Auditors verify integrity later without the data being publicly visible. Selective-disclosure credentials use a single signed credential to support disclosure of specific fields to specific parties without re-issuance. Optional zero-knowledge proofs can be deployed where a regulation requires public proof of meeting a threshold without full disclosure (for example, a requirement that emissions \leq threshold).

B.5 Deployment models

Permissionless public ledgers offer strong public verifiability and censorship-resistance, but they present confidentiality and cost-management constraints. Permissioned consortium ledgers offer stronger privacy controls and can be used for integrated supply chains. Hybrid deployments place operational data on a permissioned ledger

and anchor periodic commitments to a public ledger for stronger immutability guarantees. The architecture takes no position on which deployment is appropriate for any given sector or jurisdiction, as that is a policy choice.

Appendix C. Transfer-evidence adapter

The transfer-linked movement invariant in §3.2 requires that emissions move with the legally relevant transfer of the underlying product. When that transfer occurs and what evidentiary form establishes it are jurisdiction specific. For instance, India's GST regime treats a registered e-invoice as the canonical record of supply. The European Union relies on customs declarations and VAT records for cross-border movement. The United States relies primarily on commercial invoices supplemented by sector-specific filings. Many jurisdictions use stamps, statutory fees, tax assessments, and other fiscal levies as legally significant markers.

Because no single transfer-evidence document is universal, the architecture introduces a transfer-evidence adapter. This is a normalization layer that maps jurisdiction-specific transfer-evidence records into a common machine-readable structure for the E-ledger. The adapter does not flatten meaning across jurisdictions. It preserves it. Each record retains its own evidentiary weight under the law that issued it. What the adapter standardizes is the structure: how the record is represented, addressed, and pointed to, so that downstream systems can consume the proof point without being programmed for every jurisdiction's regime.

The adapter also supports ERP and platform-level integration. Most enterprise systems can produce or ingest the adapter format with a thin connector layer, since the underlying business events that create the transfer-evidence record already exist in the source system. The architecture is therefore additive to the enterprise stack rather than substitutive.

References

Carstens, A., & Nilekani, N. (2024). Finternet: the financial system for the future. BIS Working Papers, No. 1178.

Kaplan, R. S., & Ramanna, K. (2021). Accounting for climate change. *Harvard Business Review*, 99(6), 120-131.

Kaplan, R. S., & Ramanna, K. (2025). E-ledgers carbon accounting. SSRN Working Paper 5376839.

Kaplan, R. S., Ramanna, K., Roston, M. (2023). Accounting for carbon offsets. *Harvard Business Review* 101(4), 126-137.

Ramanna, K., Holloway, L., Panjkov, L., Zhang, C. W., & Kaplan, R. S. (2025). A proto-standard for carbon accounting and auditing using the E-liability method (v. 2). SSRN Working Paper 5427956.

Ramanna, K., Roston, S., Holloway, L., & Kaplan, R. S. (2026). A proto-standard for carbon removals accounting and auditing using the E-asset method. Working Paper.

Tables and figures

Table 1, Summary of invariants

Invariant	Architectural implication
Counted-once recording	Downstream entities receive upstream emissions as embedded emissions rather than re-originating them
Transfer-linked movement	Emissions allocation sits on the product token and travels in the same transaction logic as the product
Full allocation under causal logic	Allocation method is explicit and version-tracked; unallocated overflow is captured and flagged automatically
Conservation (incl. gross-history retention)	Dependency links and allocation controls identify possible under-and over-consumption
Append-only correction	Corrections add new attested records rather than rewriting prior ones
Layered disclosure	Shared layer carries only what is needed for integrity; fuller evidence remains controlled and private
Interoperability	Schema and exchange logic travel with the record

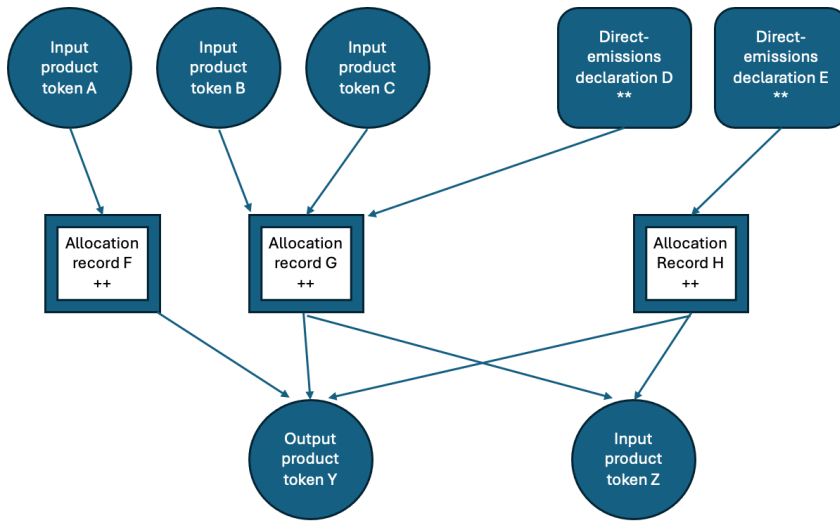
Table 2, Summary of layers of record

Layer	Primary function	Typical contents
E-ledger (entity-level)	Entity's book of account for emissions, derived from operational events	Direct-emissions records, allocations, transferred-in positions, transferred-out positions, inventory positions, attestation references
Shared layer	Cross-organizational coordination of custody, transfer integrity, and attestation references	Product token identifiers, state transitions, dependency references, attestation references, permitted shared metadata
E-ledger statement	Derived view for reporting, prepared according to a given disclosure standard	Opening and closing balances, direct emissions, emissions transferred in and out, E-assets, netting events, impairments, clarifying footnotes

Table 3, Stylized E-ledger statement for FabricatorCo

FabricatorCo: simplified period statement excerpt	kg CO ₂
Opening gross E-liabilities	0
Emissions transferred in from acquired inputs	80
Direct emissions incurred	20
Emissions transferred out with sale	(50)
Closing gross E-liabilities	50

Figure 1, The three core objects at a single supply-chain node



**** Requires Direct-emissions attestation**

**** Requires Allocation method attestation**

The figure represents a stylized case with three input product tokens, two emissions assertions, and the allocation records connecting them to two output product tokens.