Implementing E-Ledger Carbon Traceability Across Varied Manufacturing Environments

Matthew Lawrenson
Tolam Earth, Inc.

November 2025

Working Paper T04-01: Initial Draft

This working paper is a draft circulated to stimulate discussion. Matthew Lawrenson is an employee of Tolam Earth, Inc.. The work received no external funding. The author has no formal affiliation with the E-Ledgers Institute or any other organization mentioned.

Abstract

Accurately and verifiably mapping, deciphering and transferring the rights to greenhouse gas (GHG) emissions across complex, multi-party supply chains remains one of the fundamental barriers to credible carbon accounting and market-based decarbonization. Recent work by the E-Ledgers Institute[1] has articulated a principles-based model that treats emissions as ledgered events, enabling verifiable recognition, transfer, and retirement of carbon obligations between and across economic actors. Complementary work by the EFI Foundation has demonstrated early feasibility through product-level case studies such as sustainable aviation fuel, illustrating how emissions may be recorded and propagated through discrete supply-chain stages.

Just like any ontology/framework, the translation of the system principles into actual industrial practice requires operating and coordinating around heterogeneous manufacturing processes, entrenched and bespoke enterprise IT systems, and strict confidentiality constraints. This paper proposes a Dual-Node Transparency Architecture that enables verified emissions data to be processed and managed in a digital environment to accompany their mirror physical goods, all without exposing proprietary information.

By aligning the mutual objectives of shared data accountability, common carbon accounting context, and trusted sourcing and movement of specific datasets across real operating environments, the dual-node approach can provide a viable pathway for cross-enterprise, audit-ready E-Ledger Accounting deployment, supporting the array of regulatory and governance systems needed to enable private markets to price-in decision-useful emissions data into all measures of product lifecycles.

[1] The E-Ledgers Institute is a not-for-profit organization that has developed the E-Ledgers methodology. This paper engages with those concepts at a high, framework level to explore how the principles could be operationalized across supply chains. References to “E-Ledgers” in this paper are descriptive and conceptual, and should not be interpreted as an endorsement of any specific implementation, governance body, or institutional model.

Introduction

Accurately accounting for greenhouse gas (GHG) emissions across global value chains remains one of the most significant unresolved challenges in climate action. The E-Ledgers Institute has recently advanced a principles-based framework that treats emissions as verifiable ledger entries, enabling real-time recognition, transfer, and retirement of carbon obligations across economic actors (Kaplan & Ramanna, 2025). In this emerging model, entity and jurisdiction-level ledgers can ultimately consolidate into a global “geological ledger,” recording anthropogenic and biogenic emissions and removals.

Recent work by the EFI Foundation (2025) demonstrates early practical application of these concepts through a Sustainable Aviation Fuel (SAF) case study tracing carbon data from production through transport, blending, and delivery. These efforts illustrate a fundamental reality: for technology systems employing E-Ledger Framework Models to scale, they must operate reliably across heterogeneous industrial environments, supply-chain structures, data management, and technology landscapes.

Industrial Reality and Traceability Requirements

Modern supply chains consist of diverse, independent entities, sub-entities and vendors operating entrenched systems such as enterprise resource planning (ERP), manufacturing execution systems (MES), and varied inventory and production controls. These IT systems are core to business operations, tightly integrated into internal processes, and handle sensitive data that organizations are unwilling to expose externally (for good business and legal reasons). Therefore, practical adoption and inter-organizational use of an E-Ledger approach to carbon accounting requires a stack and market model that:

- Integrates with existing systems without disruptive changes
- Maintains confidentiality while enabling audit-grade verification
- Ensures accuracy and consistency across organizational boundaries

Manufacturing Diversity and Data Integrity Challenges

Industrial production processes vary substantially, creating different implications for carbon attribution and traceability:

Process Context	Carbon-Ledger Challenge
Combinatorial (multiple inputs → one output)	Aggregate upstream data without distortion
Transformational (chemical/physical conversion)	Preserve emissions ownership lineage when physical identity changes
Disaggregative (one input → multiple outputs)	Allocate emissions transparently and consistently
Bulk / pooled flow (fungible inventory)	Apply FIFO/LIFO/weighted-average logic without losing auditability
Biogenic / regenerative cycles	Accurately record negative emissions and their retirement

These operational realities introduce complexities including:

Loss of strict 1:1 traceability: Real manufacturing rarely preserves unit-level lineage. Inputs are blended, batched, pooled, split, and recombined across time and processes. Carbon accounting must remain auditable even when physical tracking becomes statistical or allocated rather than direct.

Multiple inventory and valuation methods: Firms rely on entrenched accounting systems, for example, FIFO, LIFO, weighted average, batch-lot, perpetual inventory, and more, tightly embedded within ERP controls. Any carbon accounting system must interoperate with these methods without forcing disruptive redesigns of core financial or inventory processes.

Cross-enterprise chain-of-custody: Supply chains span numerous companies with different systems, standards, and confidentiality needs. Emissions data must move across organizational boundaries with verifiable provenance and without risk of double counting, tampering, or leakage.

Verification without exposing sensitive data: Companies must protect production methods, cost structures, and trade secrets. Carbon data must be verifiable by third parties while underlying proprietary details remain confidential.

Support for iterative data refinement: Carbon values often improve as more accurate upstream recursively flows back through the system. The design must allow backward-propagating updates without breaking auditability, increasing accuracy over time while preserving historical records.

Varied digital maturity: Some suppliers have advanced digital systems; others rely on manual processes. The solution must accommodate a broad range of IT capabilities without creating barriers to participation.

Objective of This Paper

This paper proposes a practical architecture for implementing E-Ledger principles across real-world supply chains. It addresses how verified emissions data can attach to material flows in a manner compatible with existing systems, variable manufacturing modes, and confidentiality constraints. The objective is to demonstrate how E-Ledgers can scale beyond conceptual design and pilot settings to support industrial deployment, regulatory interoperability, and competitive market adoption.

Proposed High-level Architecture

To operationalize E-Ledger principles across heterogeneous industrial environments, Tolam proposes a **Dual-Node Transparency Architecture**. The design integrates with existing Enterprise Resource Planning (ERP), Manufacturing Execution Systems (MES), procurement, and inventory platforms, without requiring material system changes or exposing proprietary data. Carbon data is represented via two complementary records that separate detailed internal calculation from externally verifiable proof.

1. Carbon Data Record (CDR) — Internal, Confidential

The CDR contains full emissions-calculation detail, including methodologies, activity data, allocation logic, and supporting evidence. It remains within the originating organization and is accessible only to approved verifiers. This record provides audit-grade substantiation without disclosing sensitive operational information to the public or supply-chain partners.

2. Verified Carbon Data Record (VCDR) & Token (VCDT) — External Proof

Following verification, a rolled-up, **summary record** is produced containing only essential metadata, for example system boundaries, verifier identity, issuance timestamp, and cryptographic reference to the CDR (i.e. the carbon accounting info). This Verified Carbon Data Record (VCDR) is digitally signed and recorded as a **Verified Carbon Data Token (VCDT)** on a distributed ledger.

The VCDT functions as a **privacy-preserving digital twin** of the verified, product-level carbon accounting entry. As materials move between enterprises, the corresponding VCDT, as a digital asset, is capable of moving in parallel, ensuring transfer of rights, as well as a full chain-of-custody record to input the CDR data into subsequent products.. This prevents double-use, and preserves integrity across combinatorial, transformational, disaggregative, pooled, and regenerative production flows. Each downstream record cryptographically inherits upstream lineage, enabling traceability without exposing internal data systems.

This architecture implements carbon accountability as physical digital asset inventory movement and delivers structured product carbon intensity datasets ensuring **verifiable, auditable, cross-enterprise traceability** - while maintaining and in fact augmenting operational confidentiality.

Mapping Carbon Data Across Different Systems

Each enterprise uses its own product identifiers, hierarchies, and inventory methods. To ensure seamless interoperability, the architecture employs **shared schemas and cross-reference mapping**:

- **CDR Creation:** Each firm generates CDRs from its internal ERP/MES data and emissions methodologies. Sensitive data never leaves the boundary.
- **Verification & Publication:** A verifier signs a standardized VCDR, which is then minted into a VCDT.

- **Automated Matching:** When components move between firms, the receiving entity automatically maps the incoming VCDT to its internal SKU/BOM record, mirroring financial and materials-ledger practices.
- **Chain-of-Custody by Design:** Every VCDT maintains hash-linked references to upstream VCDRs, creating a **tamper-evident lineage** from raw materials through finished goods.

The shared schema enforces a minimum set of standardized identifiers, product IDs, facility IDs, timestamps, verification signatures, while allowing each firm to preserve internal calculation models and confidential data structures.

In effect, each product carries a **verifiable carbon chain of custody**, enabling cross-enterprise emissions accountability with the same rigor as financial traceability.

Combinatorial Production – Assembling Multiple Components

In combinatorial production (assembly processes), multiple components or materials are combined to create a new product. Each input material arrives with its own VCDR containing the verified carbon "history" of that component. As the company assembles the finished product, it adds its own emissions: direct emissions generated during the assembly process (e.g. energy used, waste generated) and a share of any overhead emissions allocated to that production run.

The output of a combinatorial process is a new composite VCDR for the finished product that integrates all the prior components' verified emissions plus the new emissions added during assembly. For example, if a pen is assembled from a plastic body, a metal nib, ink, and a label; the pen's VCDR will incorporate the VCDRs of the body, nib, ink, and label (each of which in turn encapsulates upstream emissions for those parts) along with the verified emissions from the pen's assembly step. All the source VCDR identifiers are referenced in the finished product's record, forming a hierarchy or "digital bill of materials" for carbon.

When the product is shipped to the next company in the chain, its VCDT (token) is transferred as well, carrying the cryptographic proof of the product's entire carbon profile. This token transfers as a secure transaction that mirrors the transfer of physical ownership, ensuring the environmental liability moves in lockstep with the product. Through combinatorial steps, the E-Ledger maintains a continuous chain-of-custody: even as many pieces become one, the ledger links all input VCDRs to the output

VCDR, preserving auditability. An accredited verifier typically reviews the composite VCDR for the assembled product, confirming that all inputs and added emissions are properly accounted, before it is hashed and minted into a new VCDT. Thanks to standardized metadata and lineage tracking, the assembled product's carbon data remains valid, auditable, and seamlessly transferable to any downstream partner via the token.

Transformational Production – Process Transformation of Inputs

Transformational production refers to processes like chemical reactions, refining, or smelting, where input materials are chemically or physically transformed into a new material, losing their original identity. In these cases, you cannot trace a specific input "piece" in the output because the input is altered or consumed. However, the carbon accounting follows the transformation even if physical traceability breaks. The emissions from the input materials and the process are consolidated into the output's VCDR through data lineage tracking. Practically, this works similarly to an assembly case: the Data Services Unit (DSU) aggregates all the incoming materials' VCDRs (which contain the inputs' verified emissions) along with the process emissions to create a new VCDR for the output product.

For example, consider a process that turns raw chemicals A and B into product C. A and B each come with VCDRs quantifying their production emissions. When producing C, those records are pulled into the calculation, the emissions from the reaction or processing step are added, and a composite VCDR for C is formed and verified. This output VCDR maintains full lineage: it will reference the hashes of A's and B's VCDRs (establishing parent-to-child links) and document the allocation or stoichiometric logic used to attribute their emissions into product C.

Even though A and B are transformed and cannot be physically identified in C, the E-Ledger's metadata preserves the connection – meaning auditors or partners downstream can trace back from product C's record to the records of inputs A and B.

The VCDT minted for product C can carry a hash of this verified record and the verifier's signature, so the proof of the transformation's accuracy is locked on-chain. In summary, transformational processes leverage the E-Ledger's ability to chain records: every transformation step produces a new verified record that inherits the carbon liabilities of its inputs (via referenced VCDR lineage) and remains auditable and transferable just like in assembly. This ensures that no emissions "vanish" simply because a material changed form – the carbon accountability is continuous.

Disaggregative Production – One Input Split into Many Outputs

In disaggregative or decomposition processes, a single input is split into multiple distinct outputs. Common examples include refining crude oil into various fuels, processing a log into lumber and sawdust, or dividing a bulk metal ingot into many parts. In such cases, fractional allocation of the input's carbon data is required.

The E-Ledger addresses this by fractionalizing the parent VCDR into child VCDRs for each output. First, the entire input batch has a verified VCDR (the "parent" record) representing, say, 100% of that batch's emissions. When the batch is divided into pieces, the system allocates a proportional share of the parent's verified emissions to each output and creates a new VCDR for each piece. Each child VCDR carries metadata linking it back to the parent's unique ID (a cryptographic parent–child hash relationship) and notes the allocation method used (e.g. by mass, volume, energy content, or market value). This ensures transparency in how the split was effected.

For example, if a 1-ton steel billet with 2,000 kg CO_2e of embedded emissions is cut into 1,000 steel pieces, the DSU can generate 1,000 child VCDRs, each assigned 2 kg CO_2e (assuming a simple mass-based split). All these child records would reference the original billet's VCDR hash, so any verifier or partner can trace them back to the source. When a manufacturer then takes some of those pieces to make a product, the downstream aggregation happens: the system will combine the appropriate child VCDRs with the manufacturer's own emissions to form the product's VCDR , as described earlier.

Throughout disaggregation, the VCDTs ensure auditability and transferability: each output can be given its own token carrying the verified data, or a batch token can represent a collection of outputs with lineage to the parent. The key is that verification remains intact – no matter how many ways you split the original material, every piece's carbon data is cryptographically linked to the source and thus auditable. This prevents double-counting or "lost" emissions; if you add up all the child VCDRs, they will equal the parent's emissions (minus any portion allocated to waste or by-products per defined rules). The E-Ledger enforces these balances so that the carbon ledger stays consistent with material flow.

Bulk / Continuous (Flow) Production – Pooled Material Tracking

In bulk or continuous flow production, materials are processed in aggregate (e.g. large vats, continuous chemical reactors, grain elevators) without discrete batch tracking of each input. Also, manufacturers often handle bulk inventory where interchangeable inputs (like standardized commodities or parts) are mixed. This breaks the one-to-one traceability between a specific incoming VCDR and a specific output use. To handle this, the E-Ledger uses a Pooled VCDR accounting model, analogous to how financial accounting uses weighted-average cost for inventory. All incoming shipments of a homogeneous material are assigned to an inventory pool, and the system maintains a running weighted average carbon intensity for that pool.

How it works:

Pool Definition: A pool is defined for each category of material (e.g. “Stainless Steel Screws M6”) along with its unit of measure and location scope. The pool acts as a continuous carbon inventory account.

Adding to Pool: When a new batch arrives, a VCDR is created for that shipment, and its quantity and carbon value are added to the pool calculation. The DSU updates the pool’s average carbon intensity (e.g. kg CO_2e per kg of material) based on the combined total. This way, even if batches mix, the pool’s intensity reflects a blend of all verified inputs.

Withdrawal and Derived VCDR: When materials are withdrawn from the pool for use in production, the system allocates emissions proportional to the amount taken at the current pool average. It then generates a derived VCDR for the withdrawn quantity. This derived record includes the quantity, the pool’s carbon intensity at withdrawal time, and hash references to all the parent VCDRs that contributed to the pool (establishing provenance). A verifier’s digital signature on the pooling methodology may also be included to attest that the averaging method was correctly applied. The ledger logs a credit emission removal from the pool and a debit addition to the consuming process, maintaining a clear balance.

By using this pooled approach, the carbon chain-of-custody is preserved in aggregate: every unit of material withdrawn carries a proportional slice of the verified carbon history of that pool. Even though we can’t say which specific shipment a particular grain or bolt came from, we can say, “this quantity withdrawn had X kg CO_2e based on the verified average, derived from these contributing records.” The VCDR for the withdrawal is then treated like any other, it can be transferred via a VCDT token along with the products it goes into, and it maintains links to all relevant input data. Auditors can reconstruct the pool’s composition through the hash-linked record of inputs, ensuring no double-counting or gaps in carbon data. Notably, the ELDS governance layer defines approved pooling methods and requires that mass and carbon balances tally up (emissions in = emissions out + inventory). This provides assurance that the bulk accounting remains as rigorous as individual tracking. In sum, the E-Ledger’s bulk/continuous process handling allows high-volume operations to maintain verifiable carbon integrity through governed allocation rules; every gram of material still has a traceable carbon story.

Regenerative Biological Production – Natural Carbon Uptake

Regenerative production refers to processes where natural systems (farms, forests, wetlands, etc.) absorb or store carbon as part of production. For instance, agricultural practices that build soil carbon, or a forestry operation where trees sequester CO_2 in biomass, can create a net negative or stored carbon effect.

The E-Ledger can accommodate these transactions by recording carbon removals or storage as verified data records as well. In practice, this is handled by extending the data model to Verified Carbon Asset Records (VCARs) and Tokens (VCATs), which mirror VCDRs/VCDTs but for carbon assets (sequestration) rather than liabilities. Each regenerative activity (say, one season of cover-cropping on a farm that sequesters a certain amount of CO_2) would produce a VCAR detailing the amount of carbon stored, how it was measured (using standardized methodologies like soil sampling or biomass calculations), and an independent verifier's confirmation. Once verified, this record is hashed and tokenized into a VCAT on the ledger, analogous to a VCDT, complete with the verifier's digital signature in its metadata.

These tokens representing carbon removal can flow through the supply chain similar to emission tokens. For example, if a coffee supplier practices regenerative agriculture and sequesters 100 kg of CO_2 in its soil, it could furnish a VCAR/VCAT for that removed carbon. When selling coffee beans to a roaster, the supplier might transfer a portion of that carbon credit (the VCAT) along with the physical shipment, effectively delivering a product with an embedded carbon asset (negative emission) that the roaster can claim in its balance.

The E-Ledger ensures that such transfers are recorded with the same rigor: each unit of removal is unique, verified, and linked to a specific project and time. Moreover, the system's double-entry logic can offset liabilities with assets on the Environmental General Ledger. In simpler terms, regenerative production can be integrated by pairing emissions records with removal records. The metadata structures support this by having parallel fields for sequestration (e.g. type of reservoir, duration, verification standard) that mirror those for emissions, ensuring compatibility and integrity across the ledger. Through VCDRs/VCDTs and VCARs/VCATs, an end-to-end chain-of-custody is maintained not only for carbon emitted, but also for carbon stored.

This means companies in a supply chain can reliably trace and transfer the climate benefits of regenerative practices alongside the products themselves, with full auditability.

E-Ledger Mechanics (Double-Entry for Carbon)

As emissions are passed through a value chain, each movement is the environmental equivalent of a double-entry posting in financial accounting. To enable such a system, when a physical transaction occurs (goods issue, production completion, shipment), the ERP/MES event can trigger an API call to a service that executes a VCDT transfer on a Distributed Ledger. Both parties can co-sign; the Distributed Ledger records the timestamp, transaction hash, and signatures. The system ensures postings align with physical events, preventing drift between operational and carbon ledgers.

With such a process, each transaction not only transfers carbon liability but also preserves the complete upstream chain of verified carbon data, enabling buyers to retrieve the history of all inputs through their associated VCDTs.

Illustrative internal postings (manufacturer):

1. Dr WIP / Cr Raw Materials — materials consumed (move incoming VCDTs to WIP).
2. Dr Finished Goods / Cr WIP — batch completed (new VCDR verified; mint/update VCDT).
3. Dr Transfers-Out / Cr Finished Goods — stage shipment.
4. Dr Customer Inventory / Cr Transfers-Out — delivery confirmed (VCDT to buyer).
5. Dr Retired/Offset / Cr Transfers-Out — end-of-life, taxation, offset, or retirement.

This ledger-first design enables inventory listings, carbon COGS, and full provenance reports directly from immutable token flows.

Ensuring Trust, Integrity, and Interoperability

For verified carbon data to move seamlessly across enterprises and supply chains, a shared foundation for meaning, verification, and traceability is required. This is achieved through a shared data framework

(SDF) and structured metadata rules that define how carbon information is recorded, summarized, and exchanged.

Each Verified Carbon Data Record (VCDR) includes essential identifiers, such as issuing entity, relevant product or batch, system boundary, methodology, and timestamp, allowing any party to interpret the record consistently, regardless of industry or internal IT systems. The SDF effectively acts as a common language for carbon accounting, ensuring compatibility across ERP, MES, and inventory environments.

Once verified, the VCDR is cryptographically hashed and digitally signed. A corresponding Verified Carbon Data Token (VCDT) is created, embedding the verifier's signature, cryptographic reference, and traceability fields. Any alteration to the recorded data breaks the cryptographic chain, enabling tamper-evidence and reliable audit trails.

Each token maintains lineage. When carbon is combined, split, or allocated through production, new tokens inherit cryptographic links to upstream records. This creates a transparent carbon chain-of-custody, similar in rigor to financial audit trails, ensuring no carbon record can be duplicated, double-counted, or reused.

Time-stamping, immutable data storage, and ownership logging further secure the transfer of carbon accountability across supply chains. Governance rules embedded in the SDF enforce approved calculation boundaries, allocation rules, verification standards, and controls, protecting consistency while preserving confidentiality of underlying business data.

Together, the SDF, cryptographic assurance, and encoded governance deliver secure, privacy-preserving carbon traceability across complex value chains. Every unit of CO_2e can be verified, attributed, and transferred with confidence, even across heterogeneous enterprise platforms.

An open question is which entity or entities will provide the Governance of the Shared Data Framework (SDF). As the SDF will function as a public-good layer its legitimacy depends on neutral stewardship, transparency, and broad stakeholder participation.

Conclusion

By using VCDRs and VCDTs as the common currency of carbon information, companies in a supply chain can interlink their carbon footprints much like they interlink their material flows. Different SKUs, processes, and systems cease to be barriers – the E-Ledger acts as the "connective tissue" binding them under a shared data protocol.

Whether components are combined, transformed, split, mass-produced, or even grown in a field, the verified carbon data remains attached and traceable. This cross-enterprise approach turns carbon accounting into a live, transactional process rather than an isolated reporting exercise, thereby enabling true end-to-end carbon accountability in the supply chain. The robust metadata and standards of the E-Ledger ensure that this accountability is built on a foundation of integrity and interoperability, giving all stakeholders confidence in the carbon data that underpins their sustainability efforts.

References

Robert S. Kaplan and Karthik Ramanna. 2025, 'E-ledgers Carbon Accounting'. Available from: https://www.hbs.edu/ris/Publication%20Files/26-004_2a839768-5736-4b08-97c6-bac7e72c71c7.pdf

EFI Foundation 2025 'Integrated Product and Entity Level Carbon Accounting - Putting Concepts Into Practice' Available from: https://efifoundation.org/wp-content/uploads/sites/3/2025/10/EFI-Foundation-Integrated-Carbon-Accounting-System-1.pdf

