The Plumbing Problem

How E-Liabilities and Data Quality Grading Create Carbon Accounting Infrastructure

Jan 2026: Note from the LAB

Critiques, questions comments welcomed

Nick.gogerty@carbonfinancelab.com

Two complementary systems solve the carbon data interoperability crisis. E-liabilities provides the transaction architecture. CarbonGrade provides the data quality layer. Together, they enable immediate deployment with existing data while creating market incentives for continuous improvement.

A STRATEGIC FRAMEWORK FOR SUPPLY CHAIN CARBON ACCOUNTING

I. THE ARCHITECTURE GAP

Carbon accounting faces a fundamental infrastructure problem. Data exists but cannot flow. The global economy generates millions of carbon calculations annually through life cycle assessments, environmental product declarations, regulatory filings, and corporate inventories. Yet this data remains trapped in silos, incompatible across systems, unusable for the decisions that matter.

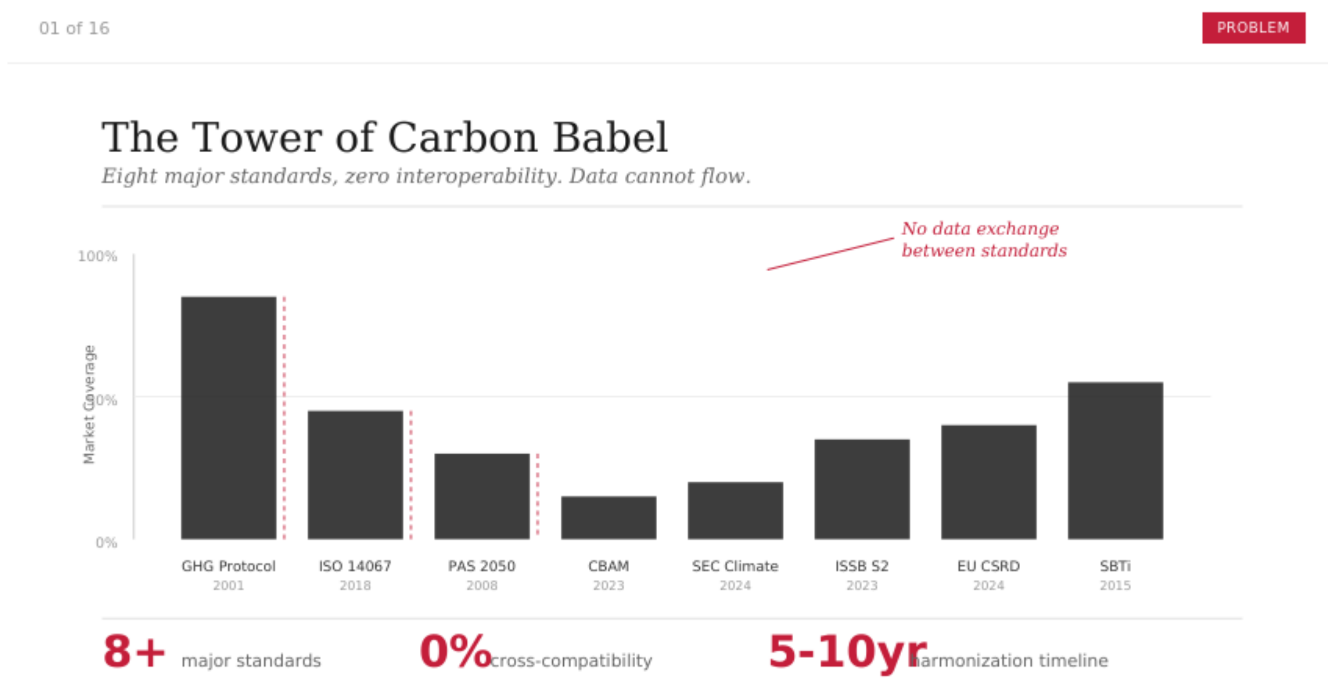

Figure 1: Eight major carbon standards with zero interoperability

The instinctive response is harmonization: create a single unified standard that everyone adopts. This approach fails. The Voluntary Carbon Market took thirty years to achieve partial methodological consensus on a comparatively simple problem. The ISO/GHG Protocol partnership targets 2028 at earliest for draft alignment, with full adoption years beyond. Climate timelines do not accommodate standards body deliberation schedules.

But the problem is more subtle than methodology fragmentation. Even if standards harmonized tomorrow, carbon data still could not flow through supply chains because the fundamental accounting architecture is wrong.

The Static vs. Dynamic Problem

Current approaches to carbon accounting are static and inventory-based. The GHG Protocol, life cycle assessments, and corporate carbon inventories all work the same way: start from a product or entity, look upstream and downstream, estimate the emissions implicated, and report a snapshot. This creates three structural problems.

First, the snapshot is already outdated when published. A product LCA takes 6-18 months to complete. By the time results are verified and formatted, the supply chain has changed. The steel supplier switched electricity sources. The logistics provider upgraded their fleet. The data describes a system that no longer exists.

Second, the same emissions get counted multiple times. When every company in a supply chain performs its own upstream estimation, the same steel plant's emissions appear in dozens of separate corporate inventories. No mechanism exists to reconcile these overlapping estimates into a coherent picture.

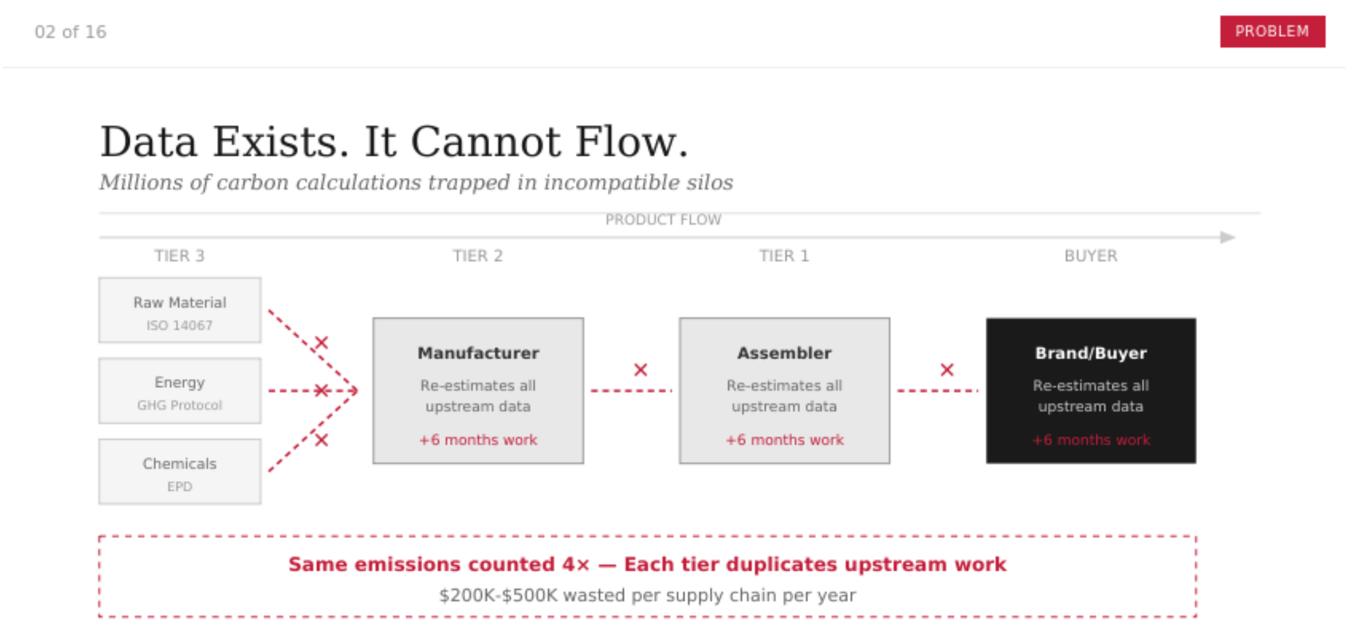

Figure 2: Same emissions counted 4× — each tier duplicates upstream estimation work

Third, the data cannot be reused. A supplier who invested in a rigorous product carbon footprint study cannot simply pass that calculation to their customers. Each customer must re-estimate the supplier's emissions using their own methodology, wasting the original investment and introducing new inconsistencies.

> *The fundamental question is not "who caused these emissions?" but "who is responsible for these emissions now?" This shift from attribution to responsibility requires transaction-based rather than inventory-based accounting.*

II. THE E-LIABILITIES SOLUTION

The E-liabilities method, developed by Professor Karthik Ramanna at Oxford's Blavatnik School of Government and the E-liability Institute, reconceives carbon accounting as a dynamic, transaction-based system. The insight is architectural: treat embedded carbon as a liability that transfers with products, exactly as financial liabilities transfer through business transactions.

Core Principles

The system operates through eight foundational principles. An entity records all direct emissions using measurement or calculation. Those emissions, verified to a reasonableness standard, are allocated to the entity's outputs using documented causal logic. When products transfer to customers through normal commercial transactions, the embedded carbon liability transfers with them.

This creates a recursive system. Each entity receives embedded emissions from suppliers on invoices, adds its own direct emissions, allocates the total to outputs, and passes the liability downstream. Carbon data flows through supply chains automatically, attached to the products and services that caused the emissions.

The VAT Parallel

The mechanism mirrors value-added tax systems. Just as VAT tracks value creation through supply chains via invoice declarations, E-liabilities track carbon creation through the same transaction infrastructure. This is not coincidental. Both systems solve the same problem: attributing cumulative upstream activity to final products in complex, dispersed supply chains.

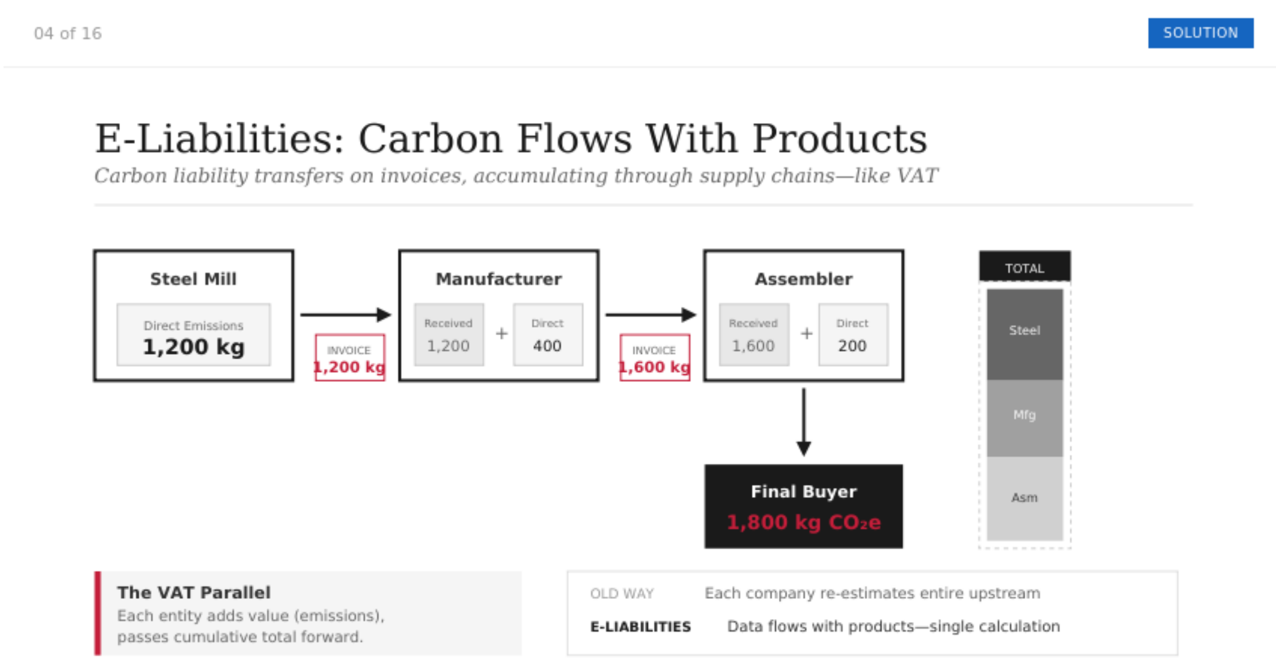

Figure 3: E-liabilities transaction flow — carbon accumulates through invoices like VAT

The VAT parallel reveals why E-liabilities can succeed where harmonization efforts fail. VAT operates across jurisdictions with different tax rates, exemptions, and administrative procedures. It does not require global uniformity. It requires only that each transaction carry forward cumulative information in a standardized format.

> *E-liabilities provides the transaction architecture. Carbon data transfers with products on invoices, accumulating through supply chains. The system works regardless of which methodology generated the upstream data.*

The Data Quality Problem

But E-liabilities introduces a challenge. If carbon liabilities transfer on invoices, what happens when upstream data varies in quality? A steel supplier using continuous emissions monitoring produces data fundamentally different from one using industry averages. Both numbers appear on invoices. How does a buyer compare them?

The E-liabilities proto-standard addresses this through Principle 5: when suppliers do not provide emissions data, purchasers record at the maximum applicable value of the product category distribution. This creates strong incentives for primary data. But it does not solve the comparability problem for data that is provided but varies in quality.

This is where the Data Quality Standard layer becomes essential.

III. THE CARBONGRADE LAYER

CarbonGrade provides the missing interoperability layer for E-liabilities implementation. Rather than standardizing methodology, it standardizes the treatment of uncertainty. The system assigns confidence grades to carbon data based on rigor and data quality, then applies conservative multipliers that make different-quality data comparable.

The Grading Mechanism

The system evaluates two dimensions. The Standard Rigor Factor assesses methodology: ISO 14067 scores higher than consultancy estimates. The Data Risk Factor assesses input quality across four dimensions: technological representativeness, temporal vintage, geographical match, and primary data share.

Combined scores map to letter grades A through F. Each grade carries a confidence multiplier applied to reported carbon values:

Grade	Multiplier	Data Type	Description
A	**1.0×**	Primary Measured	CEMS, real-time monitoring, batch-specific
B	**1.3×**	Primary Calculated	Measured inputs, verified allocation
C	**1.6×**	Supplier EPD	Supplier-specific environmental product declaration
D	**2.0×**	Industry Average	Trade association EPDs, category-level data
E	**2.5×**	Proxy/Modeled	Similar product data, estimated values
F	**3.0×**	Generic/No Data	Spend-based estimates, sector averages

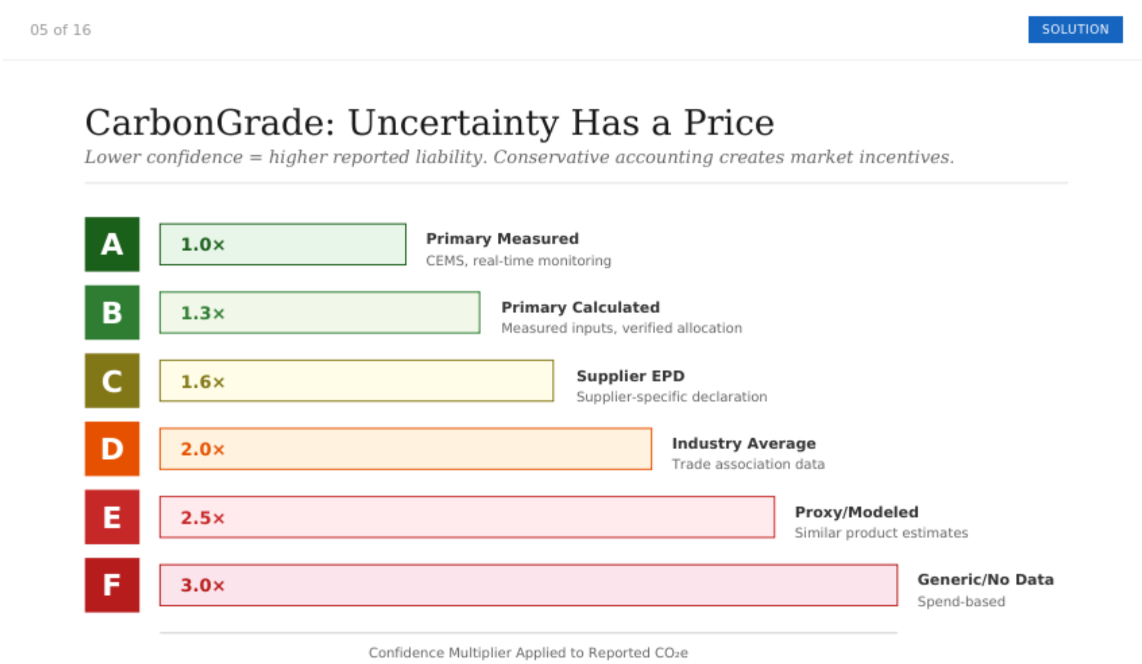

Figure 4: CarbonGrade hierarchy — lower confidence means higher multiplier

The Conservative Accounting Principle

The multipliers implement a critical insight: uncertainty must inflate reported values, not deflate them. When data quality is poor, the reported carbon liability increases. This is the opposite of how uncertainty typically gets treated, where unknown values get assigned comfortable midpoint estimates.

The logic parallels contingent liability treatment under GAAP and IFRS. When a liability is probable but uncertain in amount, accounting standards require recognizing the upper end of the plausible range. Understating liabilities misleads decision-makers. The same principle applies to carbon.

Grade D data carries a 2.0× multiplier. If a supplier reports 100 kg CO_2e using industry averages, the confidence-adjusted value is 200 kg cCO_2e. This is not a penalty. It is a conservative acknowledgment that the true value could be anywhere in the category distribution, and prudent accounting assumes the upper range.

> *The multiplier creates market incentives. Companies using poor-quality data face higher reported carbon liabilities. Investing in better measurement reduces the multiplier, directly improving competitive position without changing actual emissions.*

IV. HOW THE SYSTEMS WORK TOGETHER

E-liabilities and CarbonGrade are complementary systems solving different problems. E-liabilities provides transaction architecture: how carbon data flows through supply chains. CarbonGrade provides data quality infrastructure: how heterogeneous data becomes comparable. Neither system requires the other, but together they create something more powerful than either alone.

The Integration Point

The integration occurs at the invoice. Under E-liabilities, each transaction carries an embedded carbon liability declaration. Under CarbonGrade, that declaration includes a quality grade. The receiving entity gets two pieces of information: the reported carbon value and its confidence level.

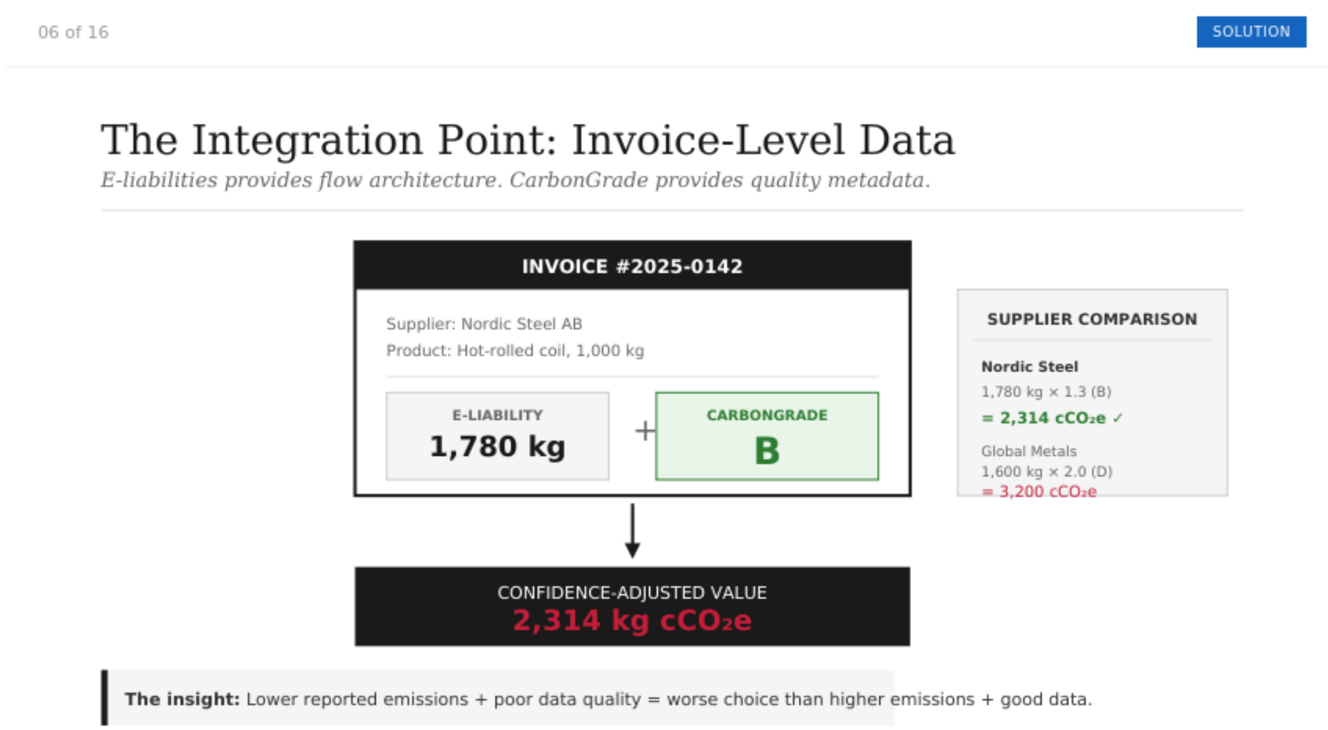

Figure 5: Invoice-level integration — emissions × grade multiplier = cCO_2e

This enables rational procurement decisions. A buyer comparing two steel suppliers sees:

Supplier	Reported CO_2e	Grade	Multiplier	cCO_2e Value
Supplier A	1,800 kg	**B**	1.3×	**2,340 kg**
Supplier B	1,600 kg	**D**	2.0×	**3,200 kg**

Supplier B reports lower emissions but has worse data quality. After confidence adjustment, Supplier A is the better choice. The buyer can rationally pay a premium for lower-carbon, higher-confidence supply.

Data Reuse Architecture

The combined system enables data reuse that neither achieves alone. Consider a furniture manufacturer receiving inputs from multiple suppliers:

The steel frame supplier provides E-liability data at Grade B, based on their verified LCA. The wood supplier provides Grade C data from their industry EPD. The fabric supplier provides Grade D data from database estimates. The hardware supplier provides no data, triggering Grade F default treatment.

The manufacturer does not re-estimate any of this data. They simply receive it through normal invoicing, apply their own direct emissions at whatever grade their measurement supports, and pass the total to customers. The grades travel with the data, enabling downstream buyers to understand cumulative uncertainty.

> *Data reuse eliminates the fundamental inefficiency of current approaches. Upstream investments in measurement quality benefit the entire downstream chain, creating positive returns on carbon data infrastructure.*

V. THE ECONOMICS OF IMPROVEMENT

The combined system creates powerful market incentives for data quality improvement. Unlike compliance-driven approaches that impose uniform requirements regardless of cost-benefit, the multiplier mechanism lets markets allocate measurement investment efficiently.

The Improvement ROI

Each grade improvement reduces the confidence multiplier, directly lowering reported carbon liability. The economics vary by product category:

Upgrade Path	cCO_2e Reduction	Typical Cost	Timeline
F → E (Database lookup)	17%	$0-1K	1-3 days
E → D (Industry EPD)	20%	$1-5K	1-2 weeks
D → C (Supplier EPD)	20%	$2-15K	2-4 weeks
C → B (Verified allocation)	19%	$5-25K	2-6 weeks
D → B (Skip C) — BEST ROI	**35%**	**$15-50K**	**4-8 weeks**

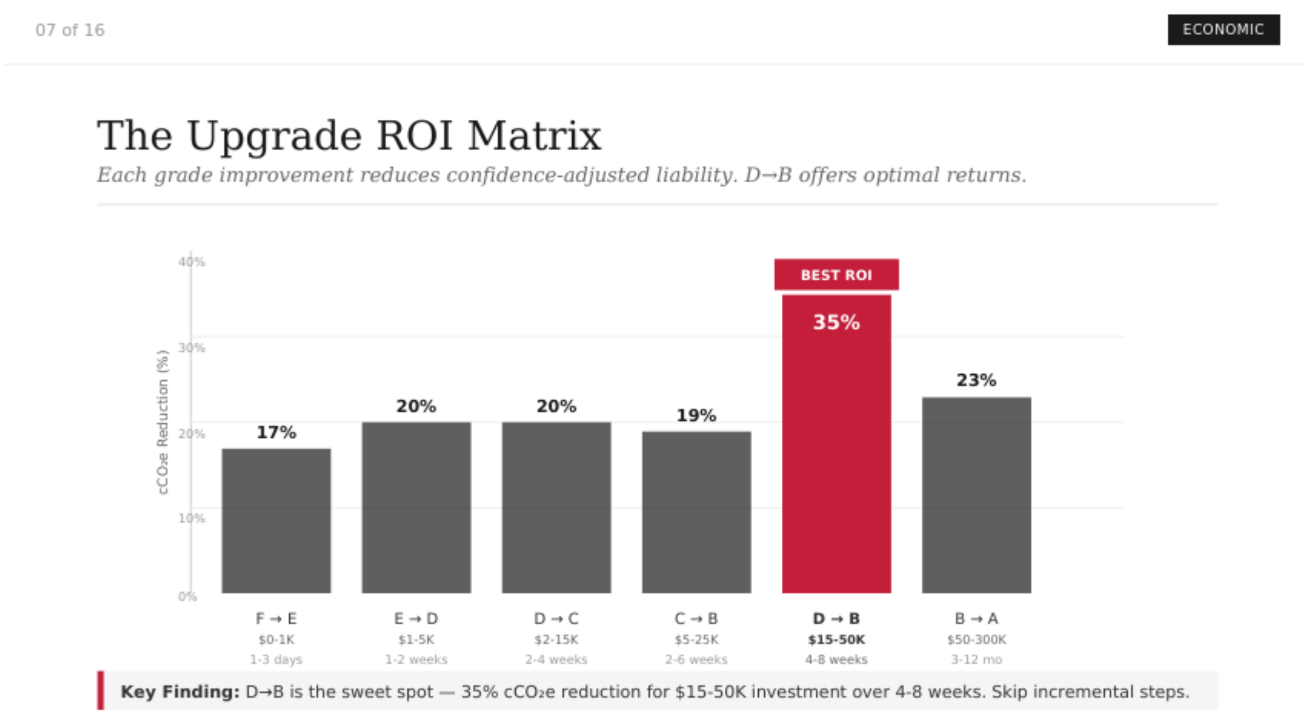

Figure 6: D→B upgrade offers optimal ROI — 35% reduction for $15-50K

The D to B upgrade path offers optimal returns for most companies: 35% reduction in confidence-adjusted carbon liability for $15-50K investment over 4-8 weeks. This is dramatically faster and cheaper than traditional LCA refreshes.

Network Effects

The system creates powerful network effects that accelerate adoption. Unlike methodology standards requiring critical mass before delivering value, the E-liabilities/CarbonGrade combination provides immediate benefits to early adopters.

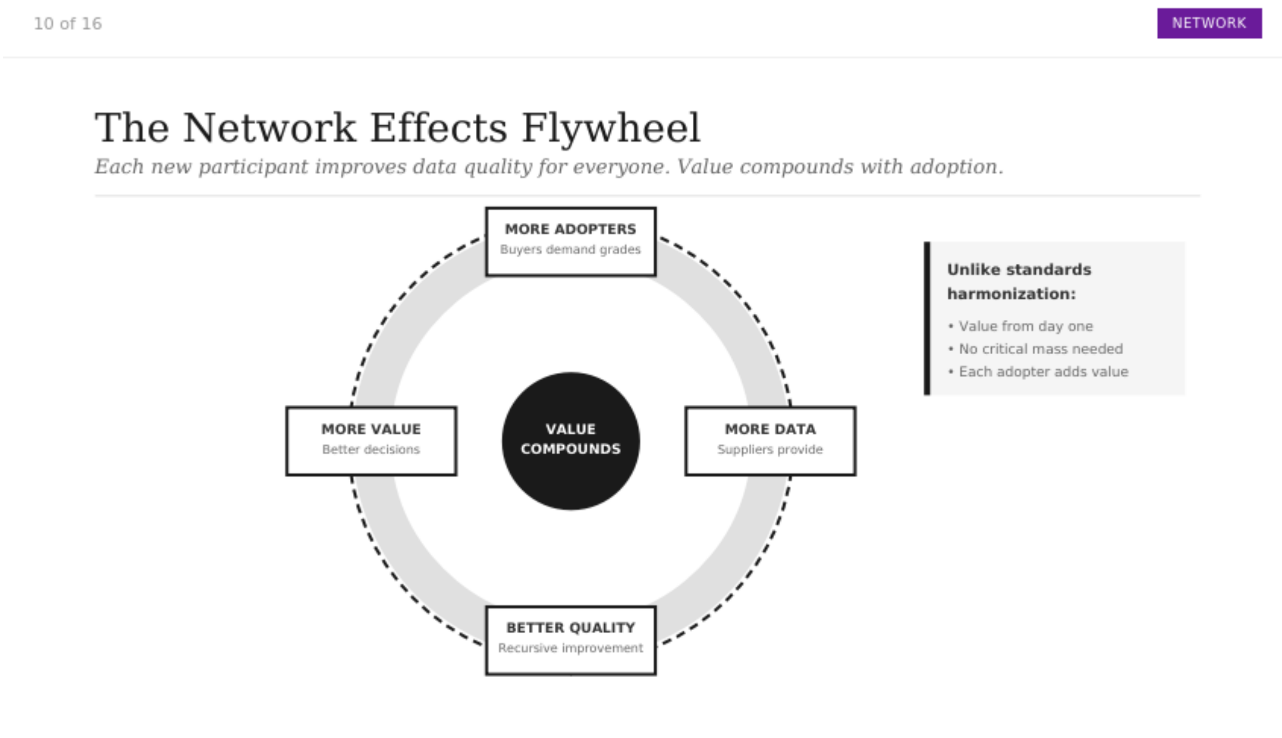

Figure 7: Network effects flywheel — value compounds with each participant

A single buyer adopting the system can immediately compare all suppliers using confidence-adjusted values. A single supplier adopting the system immediately gains competitive advantage with any buyer who interprets grades. As adoption spreads, data quality improves recursively: each entity's grade depends partly on the grades of their inputs.

VI. IMPLEMENTATION PATHWAYS

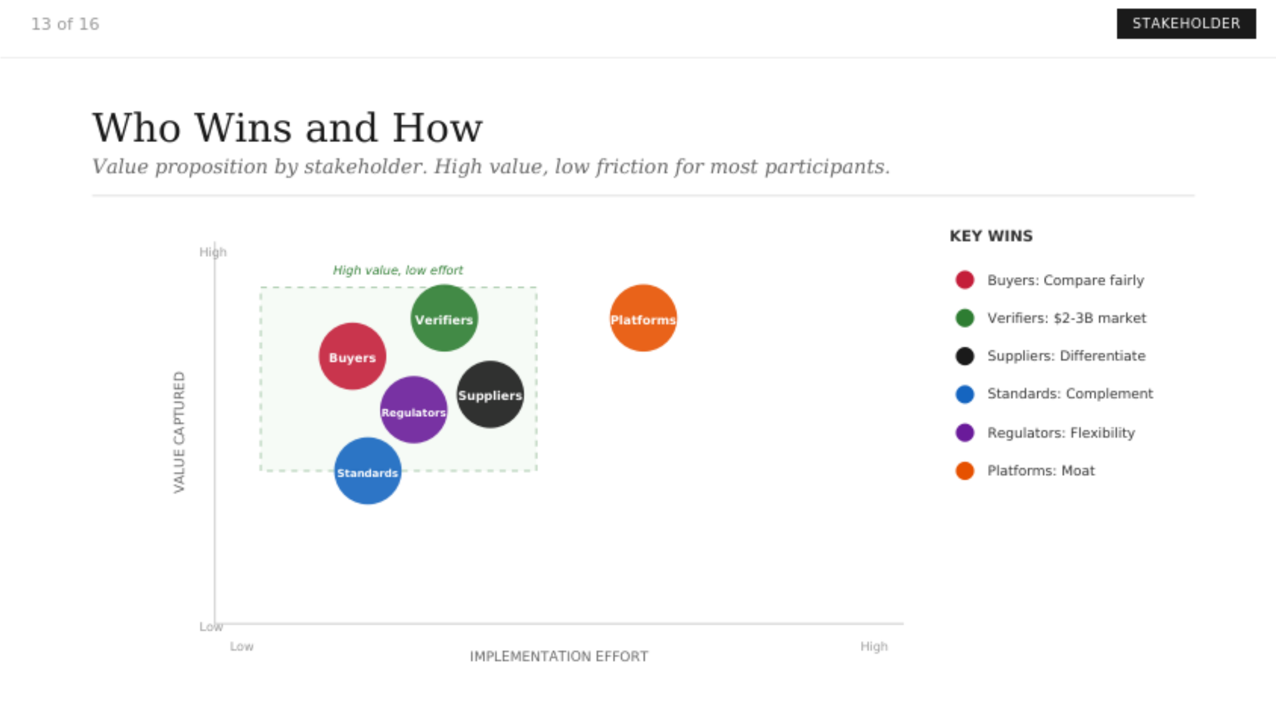

Figure 8: Stakeholder value map — high value, low effort for most participants

For Suppliers

Start with current data quality. The framework accepts any grade. A supplier with no carbon data begins at Grade F with maximum multiplier. This is not a barrier to participation. It is an honest starting point that improves with investment.

Practical first steps: Apply free databases like Ecoinvent to move from F to E. Request industry EPDs from trade associations to move from E to D. Engage with existing suppliers who have EPDs to move from D to C. These steps require minimal investment and begin demonstrating improvement trajectory to customers.

For Buyers

Use confidence-adjusted values in procurement scoring. When evaluating suppliers, compare cCO_2e rather than raw reported values. This automatically rewards suppliers who invest in data quality without requiring buyers to audit methodology.

Communicate grade requirements to supply base. A requirement like "Grade B minimum for strategic suppliers by 2027" gives suppliers clear targets and timelines for improvement investment. The framework provides the common language for these conversations.

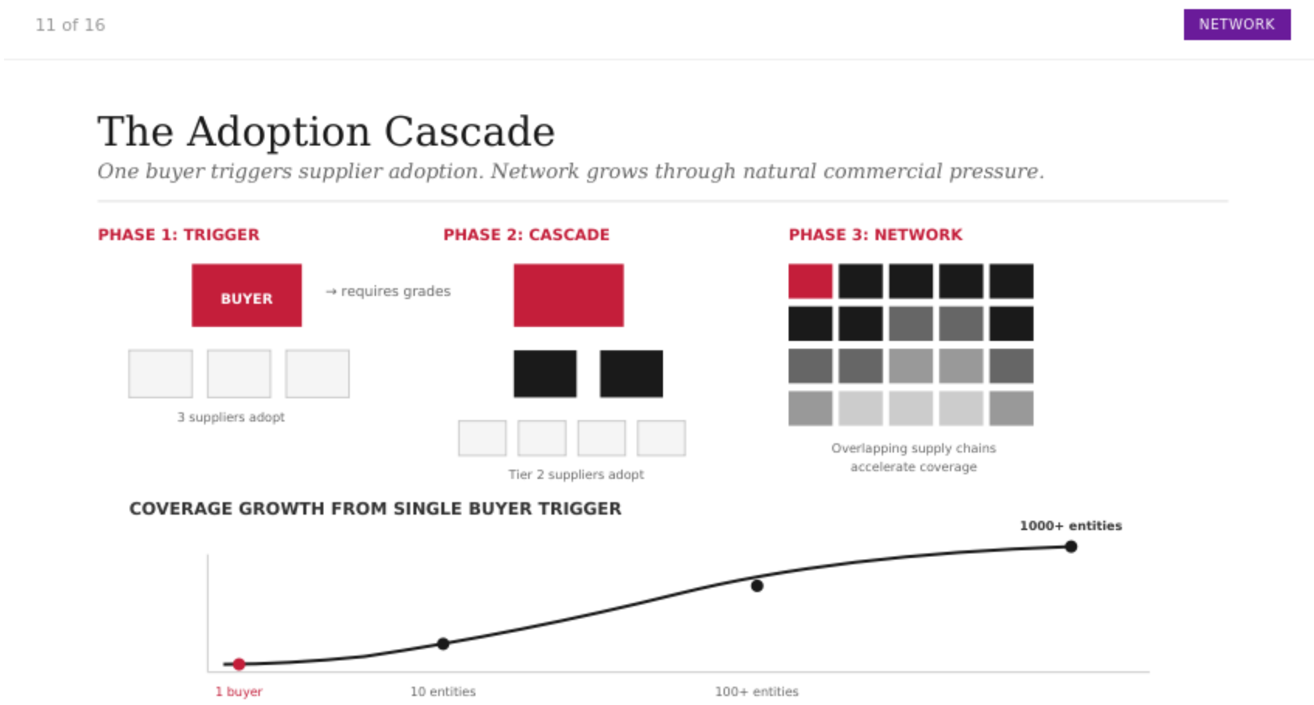

Figure 9: Adoption cascade — one buyer triggers thousands of suppliers

For Platforms and Technology Providers

Integrate translation capabilities. Procurement, ERP, and carbon accounting platforms that can ingest data from multiple methodologies and output confidence-adjusted values become essential infrastructure. First movers gain network effects as adoption spreads.

The interoperability layer is analogous to the ISIN system for securities. ISIN did not replace existing national identifiers. It created a metadata layer enabling translation between incompatible systems. Carbon data platforms should follow the same pattern.

For Standards Bodies and Regulators

The framework complements existing standards rather than replacing them. ISO 14067, GHG Protocol, CBAM, and sector-specific methodologies all continue operating. CarbonGrade simply grades their outputs, enabling comparison without requiring convergence.

Regulators can reference grades in requirements: "CBAM declarations must achieve Grade C minimum." This sets outcome standards without mandating specific methodologies, allowing compliance flexibility while ensuring data quality.

VII. THE ENDGAME

The E-liabilities/CarbonGrade combination does not preclude eventual methodological convergence. It makes convergence less urgent. Data flows immediately while standards bodies continue their work. When CBAM harmonizes with trading partners, when ISO 14067 updates complete, when GHG Protocol Scope 3 guidance matures, the framework adapts by updating grade mappings.

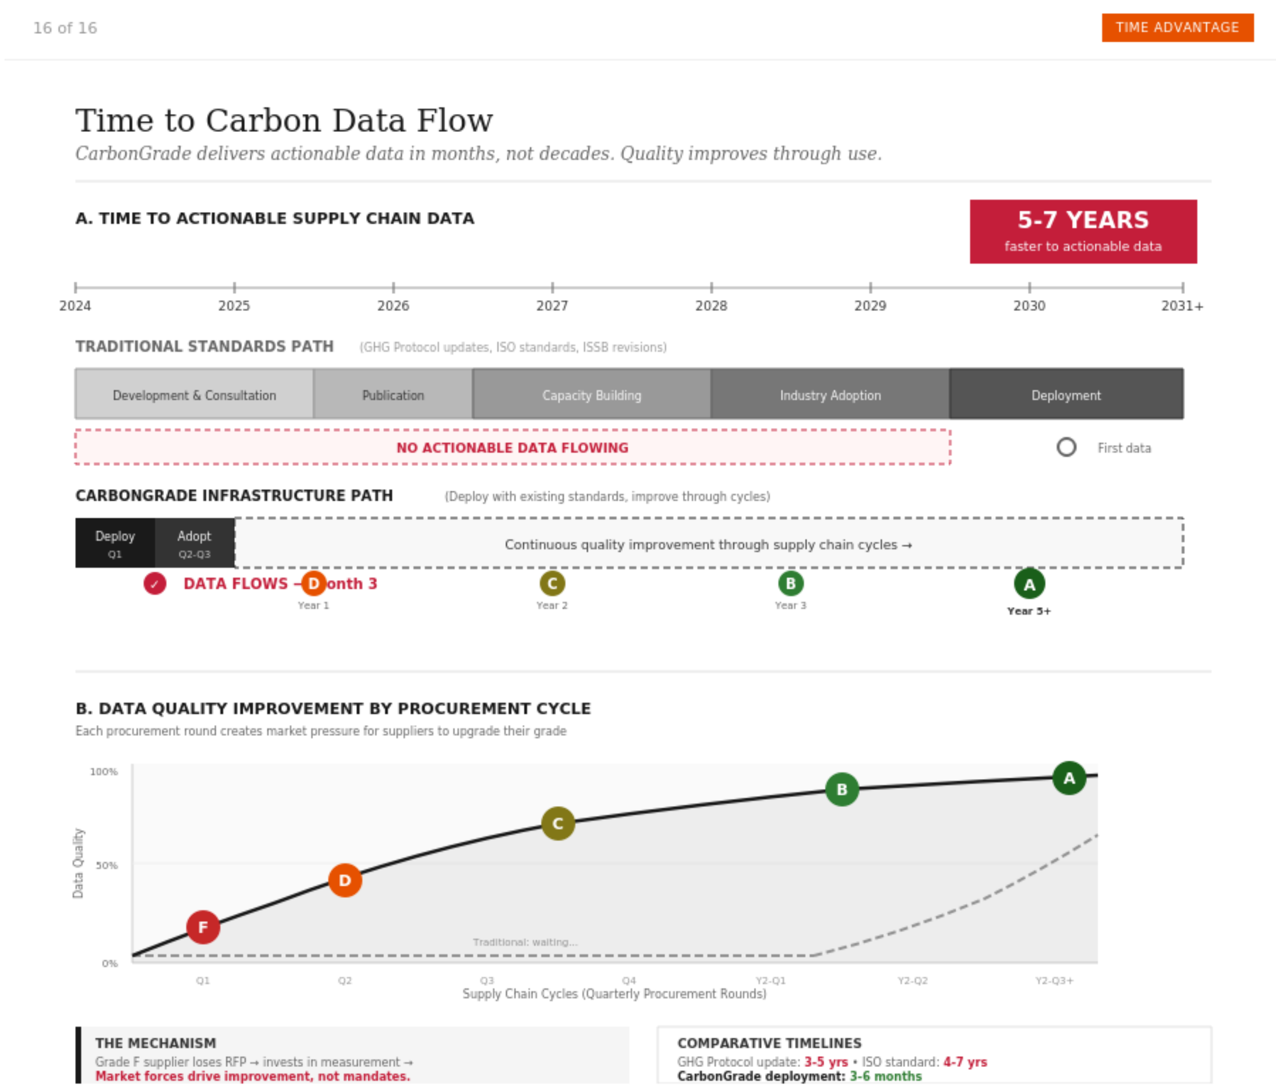

Figure 10: Time to data flow — CarbonGrade delivers actionable data 5-7 years faster

This is the key strategic insight. Waiting for perfect standards means waiting while the climate window closes. Deploying imperfect-but-improving infrastructure means action now with continuous improvement built into the architecture.

The Virtuous Cycle

The system creates a virtuous cycle across three dimensions.

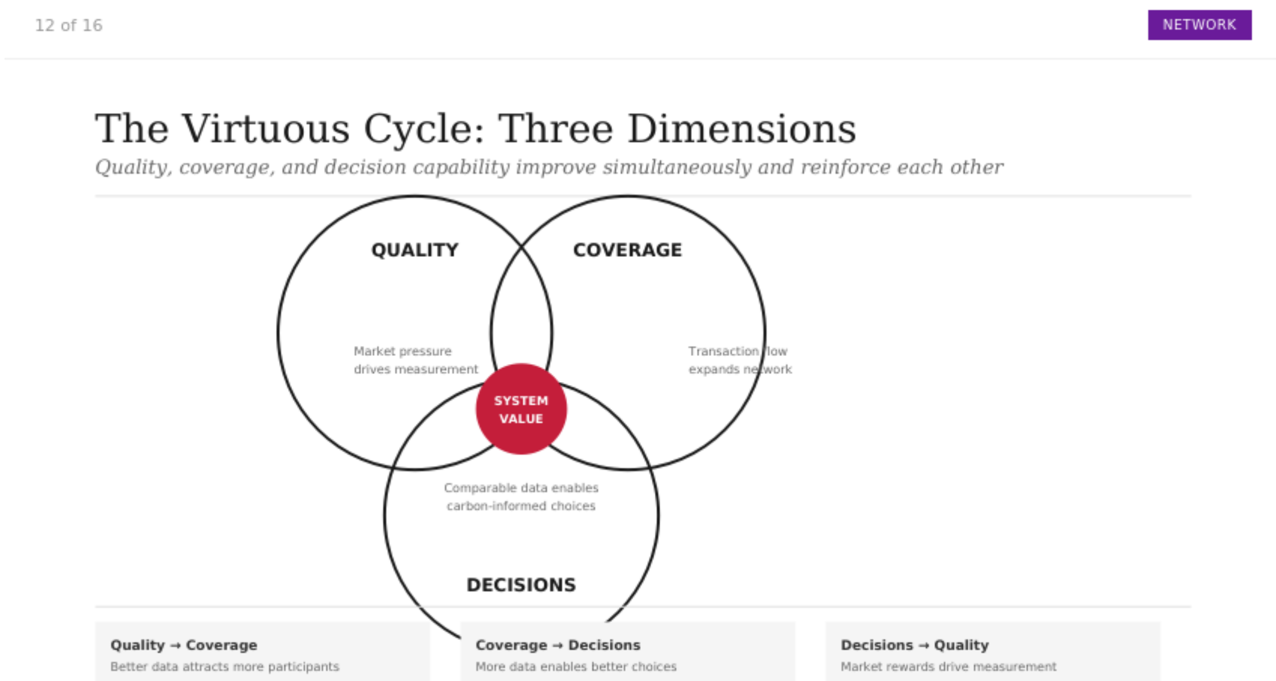

Figure 11: Quality, coverage, and decisions reinforce each other

Quality improves through market pressure. Suppliers with poor grades face competitive disadvantage. Investment in better measurement directly improves market position. Unlike compliance requirements that impose uniform costs, the multiplier mechanism lets each company optimize their investment based on competitive dynamics.

Coverage expands through transaction flow. Each new participant in the E-liabilities network adds data for their direct emissions and passes upstream data forward. Coverage grows organically through normal business transactions rather than requiring top-down mandates.

Decisions improve through comparability. As confidence-adjusted data becomes standard, procurement can systematically favor lower-carbon supply. Capital allocation can price carbon risk accurately. Product development can optimize for carbon performance. The data infrastructure enables decisions that were previously impossible.

The Principle

All accounting systems are useful fictions, models designed for purpose. The purpose today is immediate action on climate. We cannot wait for perfect standards while the window closes.

E-liabilities provides the transaction architecture. CarbonGrade provides the data quality layer. Together they create infrastructure that works with existing data, improves through market forces, and enables the decisions that drive decarbonization.

> *The goal is not a monument to accounting purity. It is functional infrastructure that gets data flowing, enables decisions, and improves through use. Accounting is the plumbing that runs empires. Build the plumbing for a decarbonized economy. Build it now.*

APPENDIX: TECHNICAL REFERENCE

E-Liabilities Core Principles

1. Record all direct emissions using measurement or calculation
2. Verify direct emissions to reasonableness standard
3. Record embedded emissions from suppliers on transfer
4. Verify transferred emissions to reasonableness standard
5. Use maximum applicable value when suppliers provide no data
6. Verify maximum value recordings to reasonableness standard
7. Allocate all emissions to current or future period outputs
8. Document allocation logic with causal reasoning

CarbonGrade Data Quality Dimensions

Dimension A: Technological Representativeness. Does data match specific production technology? Score 1.0 for exact match, 0.9 for similar technology, lower for generic category data.

Dimension B: Temporal Vintage. Data age relative to production. Under 3 years is current. 3-5 years requires no-material-change attestation. Over 5 years mandatory downgrade.

Dimension C: Geographical Representativeness. Grid and material mix match to production location. Same sub-grid scores 1.0. Continental average scores 0.5.

Dimension D: Completeness. Percentage of footprint from measured, verified inputs weighted by carbon impact. Book-and-claim or mass balance without physical segregation capped at Grade B.

Relationship Between Systems

Aspect	E-Liabilities	CarbonGrade
Purpose	Transaction-based carbon accounting	Data quality interoperability layer
Timeline	Long-term systemic change (5-10 years)	Near-term solution (6-12 months)
Adoption	Requires new accounting infrastructure	Works with existing data today
Relationship	E-liabilities achieves Grade A+	Enables E-liabilities adoption

CarbonSig Working Document