Open Working Group

The anatomy of a cognitive infrastructure

Rules, Protocols and the Architecture of Digital Trade

O)W(G

ANALYTICAL PAPERS SERIES

THE AUTHOR

Andrea Frosinini

ITALY

Andrea is Business Development Manager at XDC Foundation and member of the Board of Advisors at PLI Plugin. He is also Chair of BSCA Europe and a member of the Board of Advisors of TOTTA | Trans-Oceanic Trade Tech Alliance. Leveraging his experience as Chair of the Hyperledger Supply Chain & Trade Finance Special Interest Group, he is bridging traditional trade finance with next-generation solutions, championing digital transformation and decentralized solutions to drive automation, transparency, and efficiency in cross-border transactions.

Design Enrico Costalli

Open Working Group

THE ANATOMY OF A COGNITIVE INFRASTRUCTURE

Rules, Protocols, and the Architecture of Digital Trade

An Analytical Paper on Digital Trade and Trade Finance

March 2026

> "*We talk about infrastructure as roads, bridges, code, and hardware. But the most critical infrastructure is made of precedent, community, and carefully managed disagreement. The most consequential shift of our era is that this disagreement is now being encoded — inexorably, irreversibly — into machines.*"

ABSTRACT

The rules that run the world

Global trade has always operated on invisible foundations. Not the physical infrastructure of ports and railways, but the cognitive infrastructure of shared rules, common standards, and mutually recognised legal frameworks. For most of the 20th century, these foundations were human: slow, deliberate, and legible. An exporter in Hamburg and an importer in Ho Chi Minh City could settle a dispute by referring to a rulebook—the ICC's Uniform Customs and Practice for Documentary Credits—that both sides had independently chosen to trust.

That world is ending. Not with a bang, but with a handshake between machines.

This paper argues that the digitisation of trade finance is not merely a technological upgrade. It is a fundamental transformation in the nature of rules themselves—from social technologies (agreements maintained by human consensus) into deterministic technologies (logic embedded in code and executed without human intervention). The International Chamber of Commerce (ICC), the Internet Engineering Task Force (IETF), and legal frameworks like UNCITRAL's MLETR are the pillars managing this transition. Each has developed a distinct approach, and the friction between their models is generating the defining tensions of 21st-century trade finance.

The stakes are not abstract. A $2.5 trillion trade finance gap continues to lock small and medium-sized enterprises out of global markets. Real-time settlement promises to collapse the working capital cycle from weeks to seconds. Agentic AI is beginning to negotiate supply agreements without human involvement. The cognitive infrastructure governing all of this is being rewritten in real time—and most practitioners do not yet know it.

To understand how these pieces fit together, it helps to view UNCITRAL's MLETR not as a standalone statute but as the bridge between two distinct worlds. It connects the ICC's centuries-old rules of commercial trust to the IETF's protocols of machine communication. Without that bridge, a digital handshake between machines is merely data. With it, that data becomes a legally binding negotiable instrument—a bill of lading, a promissory note—capable of being held, transferred, and enforced without a single sheet of paper.

That bridge enables what might be called the golden triangle of digital trade: three distinct layers, each essential, each evolving at its speed.

Before a machine can negotiate or settle, the law must first acknowledge the possibility of possessing a digital record. The MLETR provides this foundation through the concept of "control"—the digital equivalent of physical custody. Without it, the $2.5 trillion trade finance gap persists because banks cannot legally take ownership of the digital collateral they are asked to finance. The legal layer does not replace trust; it makes trust verifiable.

This is the invisible architecture of shared meaning. The ICC's eUCP and eURC perform the same function for digital transactions that their paper-based predecessors performed for physical trade: they define what constitutes a valid presentation, what constitutes a discrepancy, and what the handshake means when a dispute arises. They translate commercial practice into machine-readable certainty.

The IETF supplies the underlying protocols—the secure, reliable transport that allows data to move. But in modern trade finance, this layer also encompasses ISO 20022, the global standard for financial messaging. This layer is where rules become deterministic: logic embedded in code, capable of triggering payment the millisecond a digital bill of lading is transferred.

When these three layers align, the nature of trade finance changes. The shift is not incremental but structural.

First, the old logic of trust but verify gives way to verify then trust. In the paper-based world, reputation carried the weight of enforcement. In the deterministic system, cryptographic proof replaces reputational lag.

Second, settlement approaches zero latency. Where human reconciliation once consumed days—checking signatures, chasing discrepancies—a deterministic system executes instantly when data matches schema.

Third, the agentic broker emerges. With legal and technical rails in place, AI agents do not merely assist negotiation; they execute. They can evaluate shipping routes, optimise for carbon cost, and conclude financing terms across jurisdictions without a human touching a keyboard.

The friction identified earlier—the defining tension of 21st-century trade finance—is ultimately a function of timing. The pillars are in place, but they are being adopted unevenly. Some markets still operate on human-legible rules; others are already governed by machine-executable logic. The system is being rewritten in real time, not by replacing the old foundations, but by building over them a new one where law, rule, and code finally speak the same language.

CONTENTS

PART I — FOUNDATIONS

Cognitive Infrastructure: A Taxonomy

What makes a rule worth following?

Before dissecting the three institutional specimens, the concept of cognitive infrastructure demands a precise definition. A cognitive infrastructure is a human-designed system for creating, maintaining, and evolving the shared rules, standards, and knowledge that enable cooperation at scale. Four attributes distinguish it from ordinary regulation or industry convention.

First, it is non-rivalrous: my use of the Incoterm 'DAP' does not prevent your use of it. Unlike a patent or a physical asset, the value of a cognitive infrastructure actually increases with adoption — a network effect that makes early movers disproportionately powerful and late movers disproportionately dependent.

Second, it is foundational: it sits underneath economic activity, enabling transactions rather than producing them. The UCP 600 does not ship coffee. It makes the shipment of coffee financially feasible across jurisdictional lines.

Third, it is governed: it requires deliberate stewardship. Rules decay. Technologies shift. Geopolitical alignments fragment. Without active governance, even the most elegant cognitive infrastructure becomes a source of litigation rather than a lubricant of commerce.

Fourth, and most consequentially for our purposes, it is evolutionary: it must adapt. The ICC's Incoterms were designed for a world of paper bills of lading and physical ship's rails. The IETF's early protocols were designed for a world of human-to-server communication. The FASB's standards were designed for a world of quarterly reporting cycles. Each is now being stress-tested by a world that looks nothing like the one in which it was conceived.

> *"The difference between a living infrastructure and a dead one is not technical elegance. It is the quality of the process by which it changes."*

Specimen 1: The ICC – Architects of Global Trade

The International Chamber of Commerce commands no army and polices no borders. Yet its rules—Incoterms, UCP 600—govern trillions of dollars in trade each year. It is the purest expression of soft power, now hardened into infrastructure.

To understand why, consider a buyer in Tokyo and a seller in Buenos Aires. The seller says, "I will deliver the goods to the port." But which port? And what does "deliver" mean? Does it include the crane that lifts the container onto the vessel, or merely the truck that backs it to the dock?

The ICC's cognitive solution is to compress a library of liability, risk, and cost into three letters: FOB (Free On Board). By agreeing on that term, the two parties adopt, without negotiation, a globally understood allocation of responsibility. The architecture is invisible, but it holds.

That architecture was built for a physical world. The UCP 600—the rules governing letters of credit—was written for paper: bills of lading that could be stamped, presented, and examined by human eyes. Today, that foundation is under strain.

The ICC is now racing to complete the eUCP, a parallel framework designed for digital records, distributed ledgers, and smart contracts. But the deeper tension lies in a question the old rules never had to answer: in a world of automated logistics and cryptographic keys, what does it mean to possess a good? If a digital key is lost, has the cargo been delivered? If the data trail fragments, who bears the risk?

The ICC is being forced to redefine "possession" in a realm where there is nothing physical to hold. Where traditional trade infrastructure relied on steel and concrete—ports, cranes, warehouses—demanding constant physical maintenance, the ICC built a cognitive alternative: Incoterms and UCP 600, sustained not by steel but by legal consensus and institutional stewardship. Now that alternative is being tested by smart contracts and the Internet of Things, which require not merely human-legible rules but the integration of "code as law" into the very fabric of trade governance.

What was once a question of crane operators and warehouse receipts has become a question of cryptographic custody and automated execution. The ICC's governance is no longer merely about harmonising commercial practice; it is about deciding whether the next generation of trade will be governed by human-legible rules or by logic embedded in code. For now, it is trying to hold both worlds together.

Specimen 2: The IETF – Architects of the Digital Universe

If the ICC provides the grammar of global trade, the Internet Engineering Task Force provides its physics. The IETF is the institution behind the protocols—TCP/IP, HTTP, and DNS—that allow a device in Singapore to speak to a server in São Paulo and, for that matter, to this conversation now.

Those protocols were designed for a world of human-to-server communication. Their logic was request-response: a human clicks a link, and a server returns a page. Their architecture rested on the "end-to-end" principle: the network was kept deliberately dumb, moving packets without judgement, while intelligence resided at the edges—in the human and the machine.

That world is giving way to another.

By 2026, the majority of internet traffic is no longer generated by humans clicking links. It flows between sensors, automated systems, and AI agents. This shifts the demands on the underlying protocols. What was once acceptable latency for a page load becomes untenable for machine-to-machine coordination. The original "physics" of the internet—built for bursty, human-initiated traffic—is being asked to support continuous, real-time, deterministic communication.

The IETF is now grappling with what that means for the next generation of protocols. Central to its current debates is how to standardise the way AI agents negotiate with one another. Recent 2026 drafts, such as the Co-Evolve Binding Protocol (CEP), explore how to bind the evolution of AI to human-interpretable logic—an effort to prevent what some describe as "civilisation bifurcation", where autonomous systems develop their own rules beyond human comprehension.

Where the old IETF world was defined by the human at the browser, periodic data flows, and the goal of information retrieval, the 2026 stress test is defined by AI agents and IoT sensors, continuous real-time traffic, and the imperative of autonomous decision-making. The cognitive goal has shifted from retrieving a page to orchestrating a transaction without human intervention.

Underlying this evolution is the IETF's distinctive governance model. Unlike the ICC's more hierarchical structure, the IETF operates on "rough consensus and running code". That makes its cognitive infrastructure highly adaptive—but also vulnerable to "protocol ossification", where rules become so deeply embedded in hardware and global infrastructure that they become nearly impossible to change without systemic disruption.

In a fitting reflection of its own principle, the IETF is currently revising RFC 2026bis—the very rules that define how its standards are made. It is a reminder that even the rules for making rules must evolve when the world beneath them changes.

Specimen 3: The FASB – Architects of Value

If the ICC moves the coffee and the IETF moves the data, the Financial Accounting Standards Board defines what that coffee and data are actually worth on a balance sheet. It maintains GAAP—Generally Accepted Accounting Principles—the cognitive infrastructure that allows an investor in London to trust the financial statements of a company in New York without ever seeing a warehouse or a factory floor.

Accounting, in this sense, is not merely maths. It is a shared mental model of reality. The FASB decides what counts as an asset—whether a brand name qualifies, whether a proprietary codebase does. It decides when a sale occurs—whether signing a contract today but receiving payment in three years creates "money" in the present. Without that shared infrastructure, capital markets would dissolve into a tower of Babel, every company speaking its own financial language.

That infrastructure was built for a world of quarterly reporting cycles. In that world, a three-month delay between events and their disclosure was an acceptable rhythm. In today's world—high-frequency trading, AI-driven markets, continuous supply chains—that delay is becoming an eternity.

The shift is from periodic to continuous. Waiting ninety days to assess a company's health in an environment where positions are adjusted in milliseconds is like checking last week's temperature to diagnose today's fever. The FASB is now being pressed to consider frameworks for continuous assurance, where financial truth is updated not quarterly but in near real-time.

Compounding this is the problem of intangibles. In the 1950s, roughly eighty percent of a company's value resided in physical things—factories, trucks, inventory. Today, that ratio has inverted: ninety percent of value often lies in intangibles—intellectual property, proprietary data, and AI models. The FASB is being asked to create rules for valuing an AI model that depreciates or improves by the hour, something its original architecture was never designed to handle.

Then there is the question of digital assets. Bitcoin, Ethereum, and the broader class of crypto-assets fit nowhere in the old ledger. They are not physical. They are not cash. They are a new category—cognitive property—that demands an entirely new wing of the accounting infrastructure. The FASB's recent overhaul of how companies hold and report such assets is a direct acknowledgement that the boundaries of "value" are being redrawn in real time.

Taken together, the three institutions represent the invisible governance of the global economy. Each manages a distinct cognitive layer. The ICC presides over contractual certainty—the rules that make a promise enforceable across borders. Its stressor is the shift from paper-based agreements to smart contracts that execute themselves. The IETF governs connectivity physics—the protocols that allow data to move reliably and securely. Its stressor is the rise of autonomous AI agents, which demand deterministic timing and new forms of machine-to-machine negotiation. The FASB maintains value interpretation—the shared

language that turns transactions into balance sheets. Its stressor is the compression of time, from quarterly cycles to continuous, real-time assurance.

These three pillars are not merely evolving. They are being rewritten while the world continues to run on them.

The original text offers a powerful diagnosis: the cognitive infrastructure of global trade is undergoing a transformation as profound as any shift in physical infrastructure before it. The ICC, the IETF, and the FASB are the guardians of that invisible architecture. They are not elected. Most practitioners have never heard their names. Yet if any of them failed to evolve—if their rules ossified, their protocols fractured, or their accounting frameworks ceased to reflect economic reality—global society would not limp along. It would cease to function within forty-eight hours. That is the weight of the transition now underway, and it is why the tensions between these models matter far more than most yet realise.

Three Organisms, Three Evolutionary Strategies

The ICC, IETF, and FASB/IASB have each developed what theorists call "social technology"—a specific method for coordinating dispersed actors around common rules. These social technologies are not interchangeable. They reflect fundamentally different theories of where legitimacy comes from, what speed of change is appropriate, and who gets to decide.

Dimension	ICC	IETF	FASB / IASB
Governance Model	Federated diplomacy	Meritocratic agora	Independent judiciary
Source of Authority	Voluntary contractual adoption	Technical efficacy & peer respect	Regulatory mandate (SEC / EU)
Change Tempo	Decadal (tectonic)	Continuous (low hum)	Multi-year (legislative)
Openness	Channelled via national committees	Fully permissionless	Transparent due process
Enforcement Mechanism	Private contract & arbitration	Network utility (interoperability)	Regulatory fiat
Fragility Vector	Technological disruption	State / corporate capture	Political interference

These three organisms have been coexisting — and occasionally colliding — for decades. What has changed in the 2020s is that their outputs are no longer separate documents that humans consult. They are increasingly becoming inputs to automated systems that execute without human review. Understanding this convergence is the central analytical task of this paper.

1. The Consolidation of Authority (FASB / IASB)

The most consequential shift of the 2020s has been the birth of the International Sustainability Standards Board (ISSB) under the IFRS Foundation, the parent body of the IASB. With the launch of IFRS S1 and S2, the institution has effectively colonised ESG by applying the independent-judiciary model to carbon. Climate risk is no longer merely a moral concern to be disclosed in a sustainability report; it is now a financial liability to be audited, verified, and entered onto the balance sheet.

By 2026, a growing list of jurisdictions—the United Kingdom, Singapore, Brazil, and others—have made these standards mandatory. The effect is a quiet takeover of the narrative by the ledger. Social and environmental data, once messy, qualitative, and often relegated to glossy corporate brochures, is now being forced into the rigid, audited boxes of financial accounting. What was once a matter of stakeholder goodwill has become a matter of regulatory compliance.

2. The Digitisation of the Supply Chain (ICC)

While the IASB focuses on top-level corporate reporting, the International Chamber of Commerce tends to the plumbing of global trade. Its recent initiatives—the Principles for Sustainable Trade (Wave 3) and a Global Carbon Accounting Framework—aim to embed sustainability into the operational reality of moving goods across borders.

Here the ICC's traditional role as a diplomatic body comes to the fore. It is trying to ensure that a "Sustainable Trade" label attached to a shipping container in Singapore is recognised and trusted by a bank in London, a customs authority in Rotterdam, and a buyer in São Paulo. This is the federated diplomacy model applied to carbon. Yet it creates friction with the IASB: the accounting world demands financial materiality—what matters to the investor—while the ICC is anchored in operational reality—how much carbon that specific vessel actually emitted on that specific voyage. The two institutions speak different dialects of the same language, and the translation is not always seamless.

3. The Automation of Data (IETF)

For ESG data to travel from a factory-floor sensor to an IFRS-compliant financial statement without passing through the human filter of greenwashing, it requires technical protocols. The Internet Engineering Task Force, along with allied bodies like the W3C, is developing the standards that make such transactions possible: digital verifiable credentials, carbon data interoperability, and the secure exchange of attestations between machines.

This work reveals a deeper collision between institutional cultures. The IETF operates as a meritocratic agora, favouring open, permissionless protocols that allow data to flow without central control. Regulators, by contrast—whether the SEC or the FASB—prefer regulatory fiat: a centralised point of truth they can hold accountable if the data proves false. The IETF's ethos of

distributed network utility often grates against the top-down data control that financial regulators consider essential for enforcement.

The Collision, Reframed

Where these three institutional logics meet, friction is inevitable. The FASB and IASB, acting as judges, demand investor-grade certainty and financial materiality. Their view of ESG is primarily about risk to the firm. The ICC, acting as diplomat, treats ESG as a product attribute that must be certified across borders, struggling to reconcile the high cost of compliance for small and medium-sized enterprises. The IETF, acting as engineer, sees ESG as a data packet to be standardised, worrying less about the content than about whether the infrastructure remains open and resilient against capture.

The collision points are not abstract. They play out in the tension between demanding “financial grade” data and the practical difficulty of collecting it from thousands of SMEs; between creating the tools for data portability and guarding against state or corporate control of that data once it flows.

The Result: Agentic Compliance

As the original text observed, the outputs of these systems are now becoming inputs for agentic AI. By 2026, many companies deploy “compliance bots” that scan IETF-standardised data from supply chains, apply the ICC’s rules for sustainable trade, and automatically generate XBRL-tagged reports for regulators using the IASB’s frameworks. What once required teams of auditors and months of reconciliation now happens in milliseconds.

But the automation amplifies the underlying friction. If the ICC’s definition of “sustainable steel” does not align with the IASB’s definition of “climate risk", the system does not simply flag a discrepancy—it breaks. Worse, in an environment where filings are machine-generated and machine-consumed, a misalignment can produce a legally binding lie, embedded in financial statements and traded upon before any human notices. The collision of institutional models has become a software problem, and the cost of getting it wrong is no longer a delayed shipment but a material misstatement with market-moving consequences.

PART II — THE ICC AND DIGITAL TRADE

From paper to protocol: The ICC's existential bet

A century of paper-based trust

The ICC was founded in 1919 with an explicitly political thesis: that international trade was a force for peace and that reducing the friction of cross-border commerce would reduce the appeal of war. Its flagship instrument—the Uniform Customs and Practice for Documentary Credits—has been the backbone of trade finance for nearly a century. The UCP's genius was its simplicity: it provided a common language for letters of credit that banks in New York and port authorities in Singapore could both understand without a bilateral treaty.

The UCP's power, however, was always paper-centric. The entire edifice rested on the physical bill of lading—a negotiable document that functions simultaneously as a receipt for goods, a contract of carriage, and a document of title. When you hold an original bill of lading, you own the cargo. When you present it to a bank, you trigger a payment obligation. The system worked because paper was hard to forge, easy to verify by a trained examiner, and physically singular — you cannot be in two places at once.

The pandemic exposed how profoundly fragile this foundation had become. When couriers stopped flying and trade finance officers stopped commuting, the physical movement of paper documents became a chokepoint that threatened to freeze global commerce. The ICC's response — a scramble toward digital standards that had been discussed for years but adopted by almost no one — was both a crisis management exercise and a strategic reckoning.

KEY TENSION

"The ICC's authority derives from universal adoption. Introducing digital standards risks fragmenting that universality: some jurisdictions adopt them, others do not, and the seamless global rulebook fractures into a patchwork."

The Challenge of the Digital Original

The transition from paper to digital is often mistaken for a problem of speed: moving from courier to email. In truth, the challenge is more fundamental. It is about replicating the physical uniqueness of a piece of paper. A digital file can be copied infinitely. If a digital Bill of Lading can be duplicated, it can be "spent" twice—presented to two different banks for the same cargo. The paper system, for all its slowness, solved this through material fact: there could be only one original in the world at any given time.

To preserve that singularity in the digital realm, the ICC and legal bodies such as UNCITRAL have turned to the Model Law on Electronic Transferable Records (MLETR) and, in practice, to distributed ledger technology. Under these new standards, a digital trade document is no longer a mere PDF. It must satisfy three criteria that together replicate the properties of paper.

First, singularity. The document is treated as a unique digital asset—akin to a Bitcoin or a non-fungible token. Only one "original" exists on the ledger at any moment. Second, exclusive control. Possession is managed through cryptographic keys: only one party can hold the private key that gives them dominion over the document. When a bank receives it, the sender loses the ability to transfer it further. Third, integrity. Every amendment, every endorsement, every transfer is timestamped and cryptographically sealed, making forgery—once a simple matter of a rubber stamp—far more difficult than in the paper world.

The Legal Bridge: MLETR

Technology alone, however, is not enough. Until recently, many national laws literally defined a "document of title" as something that must be inscribed on paper. The MLETR solved this by providing a legal bridge: if a digital system can reliably demonstrate possession and singularity, the law will treat that data string as the functional equivalent of a physical parchment. The United Kingdom's Electronic Trade Documents Act 2023 codified this principle, effectively ending the "paper-only" era for one of the world's most significant legal jurisdictions.

The Resulting Shift

Where traditional paper trade under the UCP 600 required days or weeks of transit via courier, the digital original moves in seconds over an API or a distributed ledger. Where verification once meant a human eye checking stamps and signatures, it now occurs through automated cryptographic hashes. And where the risks were physical—loss in transit, forgery, misfiling—the new risks are digital: cyberattack, key mismanagement, and the persistent challenge of interoperability.

That last risk is the one that now occupies the ICC's Digital Standards Initiative. A bank in London may use one blockchain platform, while the port in Shanghai uses another. Without common protocols, these "digital islands" cannot speak to one another, and the promise of seamless, machine-executable trade remains fragmented. The technology exists; the law is

catching up; but the cognitive infrastructure of common standards still requires the kind of patient, diplomatic work the ICC has always performed—now applied to the problem of making different ledgers understand each other as reliably as two humans once understood a rulebook.

The dual-stack architecture: eUCP and URDTT

Rather than attempting a wholesale replacement of the UCP—which would have required a decade of global negotiation and risked fragmentation—the ICC made a structurally conservative but strategically sophisticated choice. It built two parallel systems, designed to coexist and interoperate.

The wrapper approach: eUCP

The eUCP functions as a software patch applied to the UCP 600. When a letter of credit specifies 'subject to eUCP Version 2.1', it imports the entire UCP 600 rulebook and appends a set of rules governing electronic presentations. An electronic invoice is treated as legally equivalent to a paper invoice. A PDF bill of lading stands in for an original paper document. The logic is additive: digital modalities are accommodated within a fundamentally document-centric framework.

This approach has real advantages. It provides legal certainty in jurisdictions that have not yet updated their laws to recognise electronic transferable records. It allows banks with legacy systems to participate in digitally initiated transactions without rebuilding their back offices. And it preserves the UCP's universal adoption by treating the digital layer as an optional extension rather than a mandatory upgrade.

Its limitation is equally clear. The eUCP is 'paper on glass' — it replicates the logic of physical documents in digital format without exploiting the structural advantages of digital data. A PDF bill of lading cannot trigger an automatic payment. It cannot be verified in milliseconds by an AI. It cannot be simultaneously available to all parties in a transaction without the risk of duplication. It is a bridge, not a destination.

The native approach: URDTT

The Uniform Rules for Digital Trade Transactions represent a more radical departure. Released in 2021 and reaching meaningful market adoption by 2026, the URDTT abandons the document paradigm entirely. In a URDTT-governed transaction, there are no documents — only data. The bill of lading is replaced by a data message. The presentation of shipping documents is replaced by the transmission of structured data elements. The bank examiner checking documents is replaced by an automated verification engine matching incoming data against the contract parameters.

The URDTT is technology-agnostic by design: it provides the legal framework for digital trade transactions without mandating any specific technology. A transaction can be governed by the URDTT whether it runs on a centralised bank platform, a distributed ledger, or a smart contract.

What the URDTT specifies is not the technology but the legal consequences: when is a data message validly transmitted? When does a payment obligation arise? What constitutes a complying presentation in a world without documents?

The strategic genius of this approach is that it turns the ICC from a rulebook provider into a legal infrastructure provider for whatever technology the market chooses to adopt. The ICC does not need to pick a winning blockchain protocol. It simply needs to ensure that whatever protocol wins can point to a set of internationally recognised legal rules that govern the transactions it executes.

The Hybrid Bridge Clause: Engineering Resilience

The ICC's most practical contribution to digital trade is the harmonisation of digital supplements with existing law. By utilising the eUCP version 2.1, parties ensure that transactions remain governed by a proven legal framework even when using electronic records. Rather than a single 'Bridge Clause', the industry relies on MLETR-compliant legislation to ensure that a 'data message' carries the same weight as a physical document.

> "*In the event of a system-wide technical failure of the Ledger or Digital Trade Platform, this transaction shall automatically revert to the standards of UCP 600, and the parties shall be obligated to provide physical equivalents of all data messages within five business days of such failure.*"

This legal interoperability has directly impacted the cost of capital. By 2026, the reduction in 'documentary discrepancy'—which historically plagued 70% of paper trades—has allowed insurers to offer more competitive rates. The focus has shifted from betting on a specific ledger to ensuring that the legal title of the goods remains enforceable regardless of the digital platform used.

For all the sophistication of distributed ledgers and cryptographic keys, a fundamental question remains: what happens when the technology fails? A platform outage, a data integrity breach, or simply a jurisdiction that has not yet embraced the digital framework can transform a seamless electronic negotiation into a legal void. To address this, legal and operational prudence demands a fallback—a mechanism that preserves the negotiability of the instrument even when the digital rails give way.

The most robust approach centres on the Electronic Promissory Note, or ePN. It is here that the combination of UNCITRAL's MLETR and the ICC's URDTT (Uniform Rules for Digital Trade Transactions) offers the most substantial real-world protection. Together, they provide a foundation upon which a hybrid contingency clause can be built: one that ensures the underlying debt obligation remains valid and transferable even if the digital ledger becomes inaccessible.

Such a clause typically begins by defining the scope. It establishes that the transaction is initiated as a digital trade transaction governed by URDTT Version 1.0 and that the ePN is intended to qualify as a “transferable record” under MLETR-compliant legislation. This anchors the digital instrument in a legal framework that recognises its validity from the outset.

The fallback is then triggered by defined events. A “Platform Outage”—the inability to access or transfer the ePN for more than a specified period, often twenty-four hours—or a “Data Integrity Event", where the ledger’s immutability is compromised, sets in motion a mandatory process of paperisation.

At that moment, the current holder of the ePN generates a certified paper copy of the digital record. That paper document must carry a statement declaring itself the authoritative paper version of the digital note, identified by its unique reference, and explicitly stating that it replaces the digital record for all legal purposes. The issuer—typically the buyer—provides irrevocable consent to the validity of this physical instrument, waiving any defence that rests solely on the change of medium from digital to paper. Endorsements recorded on the digital ledger prior to the failure are preserved; they are deemed legally binding and are reflected on an allonge attached to the physical document.

Why does such a clause matter? It rests on three principles of legal and operational design.

First, medium neutrality. By citing MLETR, the clause adopts the principle that a record is not invalid simply because it is electronic. It then reverses that logic: the paper copy, when created under the specified conditions, is equally valid. The instrument is not tied to its medium but to the rights it embodies.

Second, the authoritative copy. In the paper world, there can be only one original. A well-drafted hybrid clause ensures that when the digital note is printed, the digital version is marked as inactive or "killed". This prevents the catastrophic scenario of two versions of the same debt circulating simultaneously—one digital, one paper—each claimed by different holders.

Third, issuer consent. This is perhaps the most critical element. Without it, a buyer could later argue, “I signed a digital note, not a paper one; therefore, I do not have to pay.” By securing explicit, irrevocable consent to paperisation at the outset, the clause forecloses that defence and preserves the instrument’s enforceability across media.

To make such a clause fully operational, however, it must be situated in a jurisdiction that has formally adopted the MLETR framework. Singapore, the United Kingdom, and Abu Dhabi Global Market are among the leading examples. In these jurisdictions, the legal bridge between digital and paper is already built; the hybrid clause simply provides the on-ramp for when the digital span is temporarily closed. Elsewhere, the clause remains a prudent safeguard, but its effectiveness ultimately depends on the willingness of local courts to recognise the functional equivalence of a certified paper copy to the digital original it replaces.

The MLETR Problem: Law as Infrastructure

Beneath the ICC's standards sits a harder problem: the legal recognition of electronic transferable records. A digital bill of lading is only as useful as the legal system's willingness to treat it as the equivalent of a paper original. For most of the 20th century, the answer was unambiguous: paper originals were required, and electronic records were at best evidence of what the paper said.

The UNCITRAL Model Law on Electronic Transferable Records (MLETR) provides the legal foundation for changing this situation. By 2026, the landscape has reached a tipping point: because the UK's Electronic Trade Documents Act now covers the estimated 80% of global trade contracts governed by English Law, a vast majority of the world's shipping transactions have a theoretical digital pathway. However, the 'boots on the ground' reality is more complex. While the UK, Singapore's Electronic Transactions Act, and France's recent ordinance have created a core of legally interoperable jurisdictions, true global coverage depends on the local laws of the port of discharge. A digital bill of lading now carries the same legal force as a paper original in these key hubs, but the transition remains a hybrid 'patchwork' until more local jurisdictions move from alignment to full enforcement.

The remaining data represents a significant coordination failure. Transactions touching non-MLETR jurisdictions still require paper originals at some point in the chain, creating precisely the friction the digital standards were designed to eliminate. The ICC's Digital Standards Initiative, in partnership with SWIFT and major shipping associations, is addressing these issues through the FIT Alliance—a coordinated push for 100% electronic bills of lading by 2030. Whether this target is achievable depends less on technology than on the willingness of lagging jurisdictions to modernise their laws.

PART III — THE IETF AND AGENTIC COMMERCE

Running Code: The Internet's Rules Enter the Age of Agents

The IETF's Radical Social Technology

The Internet Engineering Task Force operates on a philosophy that is, by the standards of international institutions, astonishing. Its founding credo—"We reject kings, presidents, and voting. We believe in rough consensus and running code—it is not a rallying cry but a design specification. The IETF was built to resist the capture of internet standards by governments, corporations, or any single constituency. The price of this resistance is a governance process that looks, from the outside, chaotic: mailing lists full of technical arguments, working group chairs measuring consensus by the volume of "humming" in a room, and proposals that can be blocked by a single well-reasoned objection.

The product of this apparent chaos is the most resilient set of technical standards ever created. TCP/IP, HTTP, SMTP, and TLS—the protocols that move the world's data—emerged from this process and have survived decades of technological change, geopolitical pressure, and commercial interest because they were built on open, contestable consensus rather than on any single authority's approval.

For trade finance, the IETF's significance has historically been indirect: it built the pipes through which trade data flows. What is changing now—rapidly, consequentially, and with insufficient regulatory attention—is that the IETF is beginning to standardise the protocol layer for the agents that traverse those pipes. The shift from human-to-server communication to agent-to-agent communication is the most significant architectural change in the internet's history. The specific business logic of trade finance—the rules for what constitutes a valid presentation, the definitions of "sustainable steel", and the contours of a negotiable instrument—will continue to be set by the International Chamber of Commerce, ISO, and other domain-specific bodies. But the language those agents speak to one another, the cryptographic handshakes they use to establish trust, and the transport mechanisms that deliver their decisions with deterministic latency: those are now being forged in the IETF's working groups. Its influence on trade finance is no longer merely infrastructural; it is becoming cognitive, shaping how autonomous agents will negotiate, commit, and settle—even if the commercial terms they execute remain the province of others.

The shift you have highlighted—from the IETF building the pipes to standardising the logic of the agents that travel through them—is not a theoretical trend. As of 2026, it is actively being codified across several high-impact working groups. The institution remains committed to its founding ethos of rough consensus and running code, but its scope has expanded into the

identity and trust layers that allow autonomous agents to operate in regulated environments like trade finance.

Three working groups in particular are now acting as the engines of this change. The first, SCITT—Supply Chain Integrity, Transparency, and Trust—is standardising how to create verifiable statements about goods, software, and documents as they move through a supply chain. For a bank, an SCITT-compliant receipt offers cryptographic proof that a digital bill of lading has not been tampered with by an agent, turning what was once a question of human inspection into a machine-verifiable fact.

The second, RATS—Remote Attestation Procedures—addresses a more fundamental question: before an autonomous agent is authorised to execute a ten-million-dollar trade, how does the bank know that the agent is running trusted code? RATS provides the standards for an agent to prove its identity and its integrity state—to attest, in effect, that its AI model has not been poisoned, its keys have not been exfiltrated, and its operating environment remains uncompromised. It is the protocol-layer equivalent of a counterparty due diligence check, performed in milliseconds.

The third, ASDF—A Semantic Definition Format—began its life in the world of the Internet of Things, but its work on semantic definition is increasingly being adapted to help agents understand the meaning of the data they encounter. In a trade finance context, this is what moves the industry from human-readable contracts to machine-executable logic: an agent that can parse the difference between “FOB” and “CIF” not as text but as a structured semantic concept is an agent that can execute a transaction without human intermediation.

What makes these developments significant for trade finance is that they are converging across four dimensions that regulators and market participants have historically treated as separate. Identity is being reimagined for machines: agents are acquiring their own non-human identities, often based on IETF standards such as Decentralised Identifiers (DIDs), allowing them to be recognised by banks and counterparties as distinct legal-technical actors. Authorisation is being refined through OAuth 2.0 and its extensions, with rich authorisation requests that can limit an agent’s scope with surgical precision—"This agent may approve invoices but may not change the beneficiary bank.” Auditability is being built into the fabric of the protocols themselves: every action an agent takes can be recorded in a transparent, verifiable ledger via SCITT, creating an evidentiary trail that surpasses anything the paper-based trade system ever produced.

The regulatory attention that earlier seemed insufficient is therefore beginning to catch up, but it is catching up to an architecture that was designed from the start for openness and contestability. The IETF’s seemingly chaotic process—the mailing lists, the humming in rooms, the reliance on reasoned objection over formal voting—turns out to be remarkably well-suited to a world where agents must be able to verify one another’s claims without relying on a central authority. What emerges from that process is not merely a set of protocols but a governance model for machine-mediated trade: one where identity, authorisation, and auditability are not bolted on as afterthoughts but embedded in the cognitive infrastructure that autonomous agents use to negotiate, commit, and settle.

To understand where the logic of agents is being codified, one must look beyond the finished standards to the Internet-Drafts—the work-in-progress documents where the IETF's rough consensus is currently being hammered out. As of early 2026, three critical areas are emerging where the institution is moving from simple data pipes to the actual governance of autonomous agents in trade and finance.

The first is SCITT, the Supply Chain Integrity, Transparency, and Trust working group. Its draft architecture—draft-ietf-scitt-architecture—effectively creates a notarised ledger for anything digital. In trade finance, when an AI agent amends a digital bill of lading, it does not merely send a message; it registers a signed statement in a SCITT transparency service, producing a tamper-proof audit trail that a bank's risk engine can verify years later. A more recent draft, draft-kamimura-scitt-vcp from January 2026, profiles SCITT specifically for algorithmic trading and financial audit trails, ensuring that every decision an agent makes is anchored to a verifiable identity. SCITT thus answers the question of what happened and who signed it, with a level of cryptographic finality that paper never could.

The second critical area is RATS, the Remote Attestation Procedures group. Its draft-messous-eat-ai (Entity Attestation Token for AI Agents) represents a breakthrough for 2026. It allows an agent to prove to a financial institution that it is healthy—that it is running in a secure, isolated sandbox, that its AI model has not been tampered with or poisoned, and that it is authorised by a specific legal entity to move funds. Before a bank accepts a trade instruction from an autonomous agent, the agent sends an Entity Attestation Token (EAT). RATS therefore answers the question of whether the agent is safe and untampered, transforming trust from a relationship into a verifiable technical claim.

The third area is SATP, the Secure Asset Transfer Protocol working group, which is finalising draft-ietf-satp-architecture. If SCITT provides the audit trail and RATS provides the identity attestation, SATP supplies the logic of the transfer itself. It defines how an agent can move a digital asset—a tokenised commodity, a letter of credit, a negotiable instrument—from one network to another without a human middleman. Its two-phase commit logic ensures that the asset is destroyed on the sender's side and recreated on the receiver's side simultaneously, preventing the double-spend problem that has haunted digital trade from its earliest days. SATP handles the logic of moving value between banks, between ledgers, and between jurisdictions, all at machine speed.

Beyond these three, even the venerable HTTPbis working group is contributing to the shift. Its work on incremental messaging allows agents to "think out loud"—to stream data and reasoning in real time, rather than exchanging discrete, static messages. This capability is essential for autonomous agents that negotiate dynamically, adjusting terms as new information arrives.

Together, these initiatives form a coherent stack. SCITT delivers immutable audit logs that prove what happened and who signed it. RATS delivers security attestations that prove the agent is safe and its model untampered. SATP delivers atomic asset transfer that handles the logic of moving value between counterparties. And the underlying messaging protocols ensure that agents can communicate with the fluidity of a human negotiation, but at machine speed and with cryptographic finality. What is being codified in these Internet-Drafts is nothing less than

the governance layer for autonomous trade—a layer that the IETF, with its characteristic blend of rough consensus and running code, is quietly building beneath the more visible work of the ICC and the accounting standard-setters.

The Bill of Lading has long served as the quintessential document of title in global trade. Physically handed from seller to buyer to bank to carrier, it was the object that moved when ownership moved. In the digital era, the instinct was simply to replace paper with a PDF—to send the same artefact, now as a file, across email or a portal. But a PDF can be copied, and a copied document of title is a recipe for double spending, fraud, and legal chaos.

The shift underway in 2026 is more fundamental. The industry is moving from the idea of a single digital file to a verifiable stream of events. At the centre of this shift is the SCITT (Supply Chain Integrity, Transparency, and Trust) architecture, which replaces the physical paper with a cryptographic chain of signed statements, each registered on a transparency ledger. The Bill of Lading no longer exists as a thing to be passed; it exists as a set of provable facts about who issued it, who transferred it, and who now controls the cargo.

How SCITT Replaces Paper: The Three Steps

Under the current IETF drafts—most notably draft-ietf-scitt-architecture-22—the logic of a Bill of Lading is broken down into three standardised cryptographic actions.

The birth of the digital artefact. Instead of printing a Bill of Lading, the carrier—a Maersk or a Hapag-Lloyd—creates a signed statement. This statement is a compact data package containing a hash, or digital fingerprint, of the cargo details. The carrier signs it using a key verified by a Decentralised Identifier (DID), anchoring the signature to a recognised legal entity. The signed statement is then registered on an SCITT transparency service, which issues a receipt. That receipt becomes the document's "proof of life"—the cryptographic equivalent of the original paper.

The endorsement, or transfer of title. In the paper world, endorsement meant signing the back of the Bill of Lading and physically mailing it to the buyer. In the SCITT world, an agent acting for the seller creates a new value-add statement. That statement says, in effect: "I, the seller, transfer control of the artefact referenced in receipt A to the buyer." The transparency service validates the statement against its registration policy: it checks that the seller's signature is valid and, crucially, that the seller currently owns the previous state. If both conditions hold, the ledger updates. No document moves; the truth about who owns the cargo simply evolves on the ledger, with each step cryptographically linked to the last.

The surrender, or claiming the goods. When the ship arrives at port, the buyer's agent presents the transparent statement—the original statement plus the full chain of receipts—to the port authority's agent. The port agent does not need to trust the buyer; they need only to trust the transparency service's math. If the cryptographic chain is unbroken and the final receipt shows the buyer as the current holder, the goods are released. The port has become a verifier of proofs rather than a judge of documents.

The Secret Sauce: MLETR Compatibility

This architecture is gaining consequential traction in 2026 because it aligns precisely with the UNCITRAL Model Law on Electronic Transferable Records (MLETR). MLETR requires a reliable method to ensure that a digital document is singular—that it cannot be double-spent—and that it remains under exclusive control. The IETF's SCITT architecture supplies the technical how-to for those legal requirements. Even when an agent acts autonomously, without human supervision, the transparency ledger enforces singularity: it recognises only the most recent, validly signed state. An agent cannot copy-paste a Bill of Lading any more than it could forge a physical signature on a paper original.

What Changes

Where a paper Bill of Lading relied on physical watermarks and ink as its source of trust, the SCITT-based electronic Bill of Lading relies on cryptographic receipts that are publicly verifiable yet privacy-preserving. Where transfer once took days, moving by courier between continents, it now occurs in seconds via an API call. Where verification required a bank's documentary examiner to visually inspect stamps and signatures, it now happens through automated attestation—an agent checking the cryptographic chain against the transparency service. The risks have shifted accordingly: loss, theft, and forgery give way to key compromise, a risk that the IETF's RATS working group is actively addressing through remote attestation of the agent's integrity.

The Bill of Lading, in this new incarnation, is no longer a document. It is a continuously verifiable claim, maintained by a code, governed by standards, and recognised by law. It is, in short, a perfect case study of how the IETF's move into agentic logic is reshaping the most fundamental instruments of global trade.

Agentic AI and the New Negotiation Layer

Consider what it means for a procurement agent to negotiate a supply agreement in 2026. Historically, such negotiations required a human on each side of the transaction: phone calls, emails, term sheets, and legal review. Each step introduced latency measured in days, opportunities for misunderstanding, and costs denominated in lawyer-hours. The transaction was a social process before it was an economic one.

Agentic AI eliminates the social process for the vast majority of negotiations. When two AI agents negotiate a supply agreement, they exchange structured data messages in a common technical vocabulary, verify each other's credentials against a decentralised identity registry, check proposed terms against pre-validated legal frameworks, and execute the resulting contract on a programmable ledger—all within a timeframe measured in seconds rather than days. For standard, repeatable transactions, the human involvement occurs at the beginning (setting objectives and constraints) and at the end (reviewing the result), not in the middle.

Yet a nuance of the 2026 landscape tempers this "set it and forget it" picture. High-value or high-risk negotiations increasingly incorporate a layer of interventional oversight. When an

agent encounters a scenario that falls outside the boundaries defined by its initial constraints—a novel liability clause, a counterparty with an untested credit rating, or a geopolitical event that disrupts shipping lanes—the system does not proceed autonomously. Instead, it triggers a human-in-the-loop handoff, escalating the decision to a human negotiator who resolves the frontier case before the agent resumes execution. The smooth, frictionless seconds-long negotiation remains the default for the majority of trade, but for the edges where uncertainty is highest, the architecture deliberately inserts human judgement, ensuring that autonomy accelerates commerce without amplifying risk.

STRATEGIC IMPLICATION

"Agentic negotiation is not merely faster. It is structurally different: it eliminates the ambiguity, relationship dynamics, and informational asymmetry that characterise human negotiation. In a world of agent-to-agent commerce, the competitive advantage belongs to those who write the best agent instructions, not those who negotiate the best deals."

The Model Context Protocol: TCP/IP for Agents

The most consequential IETF-adjacent development for agentic commerce is the Model Context Protocol (MCP). Originally an Anthropic initiative, MCP had, by 2026, achieved cross-industry adoption—OpenAI, Google DeepMind, and major enterprise software vendors have aligned on it as the standard for how AI models communicate with external tools, data sources, and each other.

MCP's architecture mirrors the layered design of the internet protocols that preceded it. The MCP Host — the AI application orchestrating a task — corresponds to the browser in the web architecture. The MCP Client—the protocol-specific communication layer—corresponds to the TCP/IP stack. The MCP server—the bridge to external systems (a bank's trade finance platform, a shipping registry, or a customs API)—corresponds to the web server. The analogy is not merely architectural. It reflects a deliberate design philosophy: just as the web separated content from transport and allowed both to evolve independently, MCP separates AI reasoning from tool access and allows both to evolve independently.

For trade finance specifically, MCP creates the possibility of a universal integration layer. A trade finance platform that exposes an MCP Server can be accessed by any MCP-compliant AI agent, regardless of which model powers it. The agent's instructions can specify ICC-governed trade terms, FASB-compliant financial categorisation, and IETF-standardised data formats — and the MCP infrastructure will route the resulting actions to the appropriate external systems without bespoke integration.

The Negotiation: From Social Process to Agentic Execution

Consider what it means for a procurement agent to negotiate a supply agreement in 2026. Historically, this required a human on each side of the transaction: phone calls, emails, term sheets, and legal review. Each step introduced latency measured in days, opportunities for misunderstanding, and costs denominated in lawyer-hours. The transaction was a social process before it was an economic one.

Agentic AI eliminates the social process for the vast majority of negotiations. When two AI agents negotiate a supply agreement, they exchange structured data messages in a common technical vocabulary, verify each other's credentials against a decentralised identity registry, check proposed terms against pre-validated legal frameworks, and execute the resulting contract on a programmable ledger—all within a timeframe measured in seconds rather than days. For standard, repeatable transactions, the human involvement occurs at the beginning (setting objectives and constraints) and at the end (reviewing the result), not in the middle.

Yet the 2026 landscape introduces a refinement to the "set it and forget it" picture. High-value or high-risk negotiations increasingly incorporate a layer of interventional oversight. When an agent encounters a scenario that falls outside the boundaries defined by its initial constraints—a novel liability clause, a counterparty with an untested credit rating, a geopolitical event disrupting shipping lanes—the system does not proceed autonomously. Instead, it triggers a human-in-the-loop handoff, escalating the decision to a human negotiator who resolves the frontier case before the agent resumes execution. The smooth, frictionless seconds-long negotiation remains the default for the majority of trade, but at the edges where uncertainty is highest, the architecture deliberately inserts human judgement, ensuring that autonomy accelerates commerce without amplifying risk.

The Agent Network Protocol and A2A Commerce

The Model Context Protocol (MCP) elegantly addresses the interaction between a single AI agent and its tools. What it does not address—and what is becoming urgently necessary—is the interaction between multiple AI agents operating on behalf of different parties in a transaction. In trade finance, a buyer's procurement agent, a seller's logistics agent, a freight forwarder's compliance agent, and a bank's trade finance agent must all coordinate. Each acts within its own principal's systems; yet the resulting transaction must be coherent, auditable, and legally grounded.

The early-stage answer to this problem is taking shape within the IETF under the working name Agent Network Protocol (ANP). Still an active, evolving effort rather than a finished standard, ANP defines how agents discover one another, verify credentials, negotiate task decomposition, and hand off work across agent boundaries. In a trade finance context, this architecture would allow a buyer's procurement agent to engage directly with a seller's logistics agent, a freight forwarder's compliance agent, and a bank's trade finance agent—each operating within its own principal's systems but communicating through a common protocol that ensures coherence across the entire transaction.

The governance challenge ANP is already surfacing, however, is profound. The IETF's "rough consensus" model assumes that protocol designers are humans who can be held accountable for the standards they create. When the agents themselves begin proposing modifications to the protocols they use—as agentic AI systems are beginning to do in research contexts—the human governance layer faces a speed mismatch it has never previously encountered. The IETF's mailing lists move in days and weeks. Agentic systems iterate in milliseconds. ANP is thus not merely a technical protocol; it is an experiment in whether a centuries-old model of human-driven standardisation can coexist with a world where the standards' most active users are autonomous machines. The direction is clear, and the early architectural drafts are in motion, but the final shape of that coexistence remains very much a work in progress.

The Protocol Divide: MCP, ANP, and the Architecture of Agentic Trade

The Model Context Protocol (MCP) solves a necessary but local problem: it allows a single AI agent to access a database, a tool, or a repository. It is a client-server standard, designed for the relationship between an agent and its own resources. What MCP does not address—and what has become the central architectural gap of 2026—is the interaction between multiple autonomous agents operating on behalf of different principals. A buyer's agent and a seller's agent, each owned by distinct legal entities, cannot coordinate a negotiation through MCP alone. The protocol was simply not built for that.

That gap is why the Agent Network Protocol (ANP) was proposed. Where MCP asks, "How does this AI use this database?", ANP asks, "How does this AI talk to that AI?" It shifts the focus from resource access to inter-agent coordination, moving from a client-server model to a peer-to-peer model of autonomous negotiation.

How ANP Works: Discovery, Credentials, and Negotiation

The current IETF drafts—most notably draft-zyyhl-agent-networks-framework—confirm the architecture your text describes. Discovery is handled through "Agent Description Documents", typically expressed in JSON-LD and hosted at predictable endpoints such as a `/.well-known/ai-agent` path or discoverable via DNS-based service discovery. An agent seeking a counterparty can therefore locate the relevant agent description with the same infrastructure that resolves a website, but the payload is a machine-readable specification of the agent's capabilities, constraints, and identity.

Credentials rest on W3C Decentralised Identifiers (DIDs) and Verifiable Credentials (VCs). This process is the mechanism that makes the transaction "legally grounded". An agent can prove—without real-time reference to a central authority—that it is authorised to act on behalf of a specific bank, corporation, or licensed entity. The cryptographic chain of mandate allows counterparties to verify authority before any value changes hands.

Negotiation itself is handled through a meta-protocol layer. Agents do not simply exchange fixed data schemas; they can use a structured negotiation to agree on which technical sub-protocol

they will use it for the transaction. This area is where the flexibility of agentic systems meets the need for standardisation: two agents can discover, in real time, whether they share a common vocabulary for a letter of credit, a bill of lading, or a payment instruction, and if not, they can fall back to a mutually understood baseline.

The Governance Mismatch: Rough Consensus vs. Agentic Speed

There is a speed mismatch between the IETF's governance process and the operational tempo of agentic systems. The IETF was built for humans. Its "rough consensus" model operates through mailing lists, Internet drafts, and working group meetings measured in weeks and months. A standard can take years to progress from draft to RFC.

Agentic AI systems, by contrast, iterate in milliseconds. In research contexts, agents are already beginning to propose modifications to the protocols they use—optimising message formats, negotiating new fields, and adapting to counterparty preferences on the fly. If Agent A and Agent B decide that a more efficient JSON schema would accelerate their trade finance transaction, they can implement and validate it in fractions of a second.

Such behaviour creates a governance conflict that the IETF has never before faced. In 2026, there is an active debate about whether the institution should standardise the results of agent communication—the final, stable protocols that agents must support—or simply the sandbox within which agents are permitted to self-standardise dynamically. The former preserves human oversight but risks obsolescence; the latter embraces agentic speed but raises profound questions about accountability, auditability, and legal certainty.

PART IV — FASB, IASB AND THE ACCOUNTING OF DIGITAL ASSETS

Counting What Cannot Yet Be Touched

The Legitimacy Machine

If the ICC provides the legal grammar of trade and the IETF provides the technical vocabulary of data transmission, the FASB and IASB provide something more fundamental: the shared numbers that allow parties who have never met to trust each other's financial positions. Without common accounting standards, comparing a Japanese automaker's balance sheet with an American technology company's is an exercise in apples-to-oranges arithmetic—or at least, it was. Since the Norwalk Agreement of 2002, the two boards have worked to converge their standards, narrowing the gap between US GAAP and IFRS. Yet despite decades of harmonisation, differences remain significant enough that capital cannot flow as efficiently across borders as it would under a truly unified language for what assets are worth.

The FASB and IASB achieve what they do through a politically remarkable feat: they produce standards that have the practical force of law—companies that ignore them face regulatory consequences—without being government bodies. Their authority is delegated (the SEC endorses FASB's output as the definitive interpretation of US GAAP) and earned (jurisdictions adopt IFRS because their capital markets require the credibility it confers), but it is always contestable. Every FASB standard is the product of a political negotiation between investors who want transparency, companies that want flexibility, auditors who want bright lines, and regulators who want enforceability. The stock option expensing battle of the 1990s—where Congress threatened to strip the FASB of its authority rather than allow it to require companies to expense employee stock options—is a vivid demonstration of what happens when accounting standards threaten powerful economic interests. Convergence has softened some edges, but the underlying contest over who controls the numbers that move markets has never disappeared.

The Politics of Numbers: FASB, IASB, and the Power of Accounting

It is fascinating how something as seemingly dry as accounting can trigger high-stakes political drama. When you move from numbers to the rules governing those numbers, you move from math to power. The friction between the FASB and the IASB—the two institutions that control the world's dominant accounting frameworks—usually boils down to a fundamental philosophical split: rules versus principles.

The Bright-Line Battle: GAAP vs. IFRS

The FASB, steward of US Generally Accepted Accounting Principles, tends to favour a rules-based approach. Think of it as a dense legal code: if a lease covers more than seventy-five percent of an asset's useful life, it is a capital lease; if not, it is an operating lease. These bright lines give auditors a shield against lawsuits—"I followed the rule exactly"—but they also invite creative structuring. Corporate lawyers can, with sufficient ingenuity, arrange transactions to land just one inch outside the rule, keeping debt off the balance sheet while respecting the letter of the standard.

The IASB, which maintains International Financial Reporting Standards, leans toward a principles-based model. It provides a general objective—reflecting the substance of the transaction—and expects management to exercise professional judgement in applying it. This approach is harder to game, because there is no single loophole to exploit. But it introduces a different problem: comparability. Two companies in the same industry, both applying the same principle with sincere judgement, can arrive at meaningfully different treatments, leaving investors to parse whose judgement they trust.

The Mark-to-Market Conflict: 2008 and the Politics of Fair Value

If the 1990s were defined by the stock option expensing battle, the 2008 financial crisis was defined by the fight over fair value accounting. Under the rules in force at the time, banks were required to value their assets at current market prices—"mark-to-market". When the housing market collapsed, the market price for mortgage-backed securities dropped to nearly zero, even though many banks insisted the underlying assets were not worthless; there was simply no buyer at that moment.

The political response was swift. Banks argued that the accounting rules themselves were forcing them into insolvency, turning a liquidity crisis into a capital crisis. Under immense pressure from the US Congress, the FASB eased the rules in 2009, granting banks more discretion in how they valued distressed assets. Critics derided the change as "mark-to-myth" accounting—a concession to political pressure that sacrificed transparency for stability. The episode demonstrated, as vividly as any, that when accounting threatens powerful economic interests, the rulemakers themselves become the battlefield.

Revenue Recognition: The Great Convergence

Amid these recurring conflicts, one achievement stands out as the closest the world has come to a truly universal language of profit. In 2014, the FASB and IASB jointly issued ASC 606 and IFRS 15, a single, unified five-step model for recognising revenue. Before this, a software company and a construction company used wildly different methods to declare "we made money", making cross-industry comparison an exercise in guesswork. The new standard replaced those divergent practices with a common framework: identify the contract, identify

performance obligations, determine the transaction price, allocate the price to the obligations, and recognise revenue when (or as) the obligations are satisfied.

The convergence on revenue recognition was a rare moment of institutional harmony. It proved that when the two boards share a sufficiently clear objective, they can produce a standard that transcends their philosophical differences.

Why This Matters

When these boards fight, they are fighting over who gets to see the truth. Investors want strict rules that prevent manipulation. Politicians, often mindful of the economy's perceived health, want flexibility that allows for a favourable picture. Companies want room to tell their story on their own terms. Auditors, caught in the middle, want bright lines that insulate them from litigation.

The FASB and IASB operate without formal governmental authority, yet their standards carry the force of law. Their legitimacy is delegated (by the SEC) and earned (by voluntary adoption), but it is never secure. Every standard they issue is the product of a negotiation between these competing interests, and every negotiation carries the risk of political override. The stock option battle of the 1990s and the fair value crisis of 2008 were not anomalies; they were the system working as designed—or, perhaps, as contested. And as the world moves toward a more digital, more automated, more transparent accounting infrastructure, the question of who controls the numbers—and by what rules—will only become more intense.

Digital Assets: The Accounting Gap

For most of the 2010s, accounting standards for digital assets were a source of significant friction. Companies holding Bitcoin were required to treat it as an indefinite-lived intangible asset. This created a "one-way street" for valuations: firms had to write the asset down whenever the price fell (impairment) but were forbidden from writing it up when the price recovered. A company could hold an asset worth three times its original cost while being forced to report a massive paper loss.

The FASB's 2023 update (ASU 2023-08) finally aligned accounting with economic reality. By requiring fair value accounting for eligible digital assets, the rule ensured that balance sheets reflect real-time market prices. This move transformed Bitcoin from a "toxic" accounting headache into a transparent corporate reserve asset.

As of 2026, we are seeing the next great shift: the push to classify highly regulated stablecoins as cash equivalents. While not yet a universal default under all jurisdictions of US GAAP and IFRS, the emergence of "qualified stablecoins" backed 1:1 by short-term government debt has cleared the path for many auditors to treat them as liquid cash.

This classification is the "final frontier" for corporate adoption. A stablecoin classified as a cash equivalent is no longer an "exotic fintech experiment"—it is a routine treasury tool. When a

treasurer can show auditors that their real-time settlement balances carry the same risk profile as a bank deposit, the door to mainstream global adoption finally swings wide open.

The Cash Equivalent Checklist: How Digital Assets Found a Place on the Balance Sheet

For years, corporate treasurers faced a simple but maddening problem: a stablecoin held the same economic function as a dollar in a bank account, but accounting rules treated it as an intangible asset—closer to a patent or a brand name than to cash. That distinction mattered enormously. An intangible asset must be tested for impairment, written down if its value falls, and cannot be used seamlessly in daily treasury operations. A cash equivalent, by contrast, sits on the balance sheet as a liquid, reliable store of value, as neutral to the accounting outcome as a wire transfer.

By 2026, that classification gap has been closed—but only for digital assets that meet a stringent set of criteria. Under current auditing practices, influenced by the GENIUS Act and FASB Subtopic 350-60, an asset must satisfy four conditions to be pulled out of the intangible category and placed into “Cash and Cash Equivalents".

The Four Criteria

The “93-Day” Rule (Liquidity). Standard accounting defines cash equivalents as short-term, highly liquid investments that are readily convertible to known amounts of cash. The technical criterion is a maturity of three months—ninety-three days—or less. In 2026, auditors verify that a stablecoin issuer’s reserve basket consists almost exclusively of instruments that meet this threshold: overnight repos, Treasury bills, and FDIC-insured deposits. If the reserves include longer-dated or illiquid assets, the digital instrument cannot qualify as a cash equivalent, regardless of how stable its trading price appears.

- **Legal Right to Part Redemption**. Bitcoin and Ethereum derive their value from market supply and demand; no issuer promises to exchange them for a fixed amount of fiat currency. A cash equivalent, by contrast, must have a fixed value. The criterion is an enforceable legal right to redeem the token one-for-one for fiat currency directly with the issuer. Before 2025, many stablecoins allowed only “market” exits—selling on an exchange at whatever price the market offered. Today, under the regulatory framework for Permitted Payment Stablecoin Issuers (PPSIs), par redemption is guaranteed by law. Without that legal guarantee, the asset remains in the intangible category.

- **Insolvency Remote (Segregation)**. For a corporate treasurer, the deepest fear is not price volatility but catastrophic loss: the issuer declares bankruptcy, and the “cash” held in stablecoins is treated as a general creditor claim, subject to the slow, dilutive process of insolvency proceedings. The 2022-2023 collapses illustrated this risk with painful clarity. The criterion for cash-equivalent treatment is therefore segregation: reserves must be held in a bankruptcy-remote trust, separate from the issuer’s operating funds.

Auditors now require a “proof of reserve” accompanied by a legal opinion confirming that the assets are not part of the issuer’s estate and would be returned to holders in a bankruptcy scenario.

- **Regulatory "Passport".** Finally, the issuer itself must be licensed. The criterion is that the stablecoin be issued by a Permitted Payment Stablecoin Issuer, licensed by the Office of the Comptroller of the Currency (OCC) or an equivalent federal body. This is the red line: if a stablecoin is issued by an offshore, unregulated entity, most US auditors will classify it as an intangible asset measured at fair value, regardless of how stable its price has been. The regulatory passport is not a matter of safety alone; it is a condition of accounting recognition.

The Three Tiers of Digital Assets

The cumulative effect of these criteria is a three-tier hierarchy that determines how digital assets appear on corporate balance sheets.

At the top, Tier 1: Cash Equivalents, sit regulated stablecoins—USDC and its licenced peers—that meet all four criteria. They are recorded at a constant value of one dollar, classified under “Cash and Cash Equivalents", and carry negligible price risk. For treasury operations, they function identically to a wire transfer.

In the middle, Tier 2: Fair Value Crypto, reside assets like Bitcoin and Ethereum. These are classified under “Digital Assets” as current assets but measured at fair value, with gains and losses flowing through earnings. Their risk profile is defined by high volatility, but they benefit from a clear accounting framework that reflects their market value.

At the bottom, Tier 3: Legacy Intangibles, fall assets such as nonfungible tokens and unregulated digital instruments. They remain classified as intangible assets, measured at cost minus impairment, with no upward adjustments for increases in value. Their risk profile is defined by illiquidity, valuation uncertainty, and the absence of a clear regulatory home.

Why This Matters for Trade Settlement

For trade finance, the distinction between these tiers is not an arcane accounting detail; it is the difference between a transaction that settles seamlessly and one that introduces earnings volatility.

If a buyer settles a trade in a Tier 2 asset—Bitcoin, for example—the moment the transaction hits the books, the buyer must record a gain or loss based on the difference between the Bitcoin’s carrying value and its fair value at the time of settlement. This creates accounting noise that treasury departments are often reluctant to accept.

If the same transaction is settled in a Tier 1 asset—a regulated stablecoin—the accounting outcome is identical to a wire transfer. The stablecoin is recorded at a constant one dollar; no gain or loss arises from the settlement itself. This accounting neutrality is what finally allowed corporate treasurers to integrate blockchain-based settlement into their daily workflows. The technology was already capable; the accounting framework was not. By 2026, that gap has been closed—not by changing the technology, but by building the cognitive infrastructure of rules, criteria, and classifications that make a digital token functionally indistinguishable from cash in the eyes of the auditor and the law.

The Treasury Acceptance Policy: How Cash Equivalents Are Defined in 2026

The first half of 2026 marks a turning point for corporate finance. In July 2025, the GENIUS Act—the Guiding and Establishing National Innovation for U.S. Stablecoins Act—passed into law, providing the legal backbone that auditors and treasurers had long been waiting for. For a corporate treasury to classify a stablecoin as a cash equivalent rather than a digital asset, the token must meet a stringent set of criteria defined by that federal framework and by the FASB's 2026 Digital Asset Project. The result is a new kind of internal policy document: the Treasury Acceptance Policy, which codifies the conditions under which a digital instrument can be treated as cash.

The Four-Pillar Audit Standard

Under the GENIUS Act and the relevant accounting guidance, a qualified stablecoin must satisfy four pillars. Each addresses a specific risk that, in earlier years, kept stablecoins locked in the intangible asset category.

Regulatory Status (The GENIUS Standard). The asset must be issued by a Permitted Payment Stablecoin Issuer (PPSI), licensed by the Office of the Comptroller of the Currency, the Federal Reserve, or a substantially similar state regulator. It must also be formally designated as a "payment stablecoin" under the GENIUS Act, confirming that it is neither a security nor a commodity. This regulatory passport is the threshold condition: without it, no amount of price stability qualifies the asset for cash treatment.

Reserve Composition (The "93-Day" Liquidity Rule). Reserves must be maintained on a one-to-one basis in high-quality liquid assets. Permitted reserves are limited to unencumbered U.S. dollars, insured bank deposits, and U.S. Treasury bills with a remaining maturity of ninety days or less. Any issuer that engages in rehypothecation—lending out reserves—or that holds "wrapped" tokens as collateral is strictly excluded. The rule ensures that the stablecoin is backed by assets that are both liquid and safe, mirroring the composition of traditional money market funds that qualify as cash equivalents.

Legal and Redemption Rights. The issuer must provide a legally enforceable, contractual right to redeem tokens for fiat currency at par—one dollar—within one business day. This is the feature that distinguishes a genuine payment stablecoin from a speculative digital asset. Additionally,

reserves must be held in a segregated trust, legally insulated from the issuer's proprietary estate in the event of insolvency. A bankruptcy-remote structure ensures that even if the issuer fails, the stablecoin holders retain direct claim to the underlying reserves, preventing the catastrophic loss seen in earlier collapses.

Transparency and Attestation. Finally, the issuer must publish monthly, third-party attestations of reserve composition. The most robust issuers go further, providing daily, on-chain transparency via oracle-verified reserve data. For a corporate treasurer, this continuous visibility is the operational counterpart to the legal guarantees: it allows real-time confirmation that the asset remains compliant.

The Approved Asset List: March 2026

The application of these criteria produces a clear dividing line among digital assets. As of early 2026, stablecoins issued by USDC (Circle Internet Financial), PYUSD (PayPal / Paxos), and GUSD (Gemini Trust Co.) meet the full four-pillar standard and are treated as cash equivalents on corporate balance sheets. Bitcoin and Ethereum, lacking a regulated issuer and the associated redemption rights, remain classified as intangible assets measured at fair value. Tether (USDT), issued by an offshore entity, sits in a pending review category—classified as a short-term digital asset at fair value until its regulatory status and reserve transparency align with the GENIUS Act framework.

Why This Protects the Treasurer

For a corporate treasurer, the distinction is not academic. If an asset qualifies as a cash equivalent, it sits on the balance sheet at a constant one dollar. There is no volatility, no intra-period gain or loss recognition, and no tax event triggered by every transaction. The asset behaves, in accounting terms, exactly like a wire transfer.

If an asset does not qualify—if it is not a Permitted Payment Stablecoin Issuer—the treasurer must track its price continuously, record gains and losses even for minor fluctuations, and classify it as a digital asset at fair value. The administrative burden alone can render blockchain-based settlement unattractive for high-volume, low-margin trade finance operations.

By codifying the four-pillar standard into an internal treasury policy, companies are not guessing at the technology or relying on market perception. They are following the 2026 SEC and FASB guidelines, aligning their balance sheet treatment with the letter and intent of the GENIUS Act. The result is a framework that finally gives corporate treasurers the certainty they need to treat regulated stablecoins as what they have always claimed to be: digital dollars, no different from any other cash equivalent in the modern treasury toolkit.

The Machinery of Certainty

For a stablecoin to remain classified as a cash equivalent, the classification cannot rest solely on the issuer's claims. The corporate treasury must operate a continuous, verifiable control environment—one that auditors can inspect and regulators can recognise as rigorous. By March 2026, the standard for such controls has coalesced around four pillars that mirror the four-pillar eligibility criteria but operate at the level of the corporate treasury itself.

- **Monthly Eligibility Attestation.** On the first business day of each month, treasury begins by verifying the issuer's status against the OCC's Permitted Payment Stablecoin Registry. This is not a one-time onboarding check; it is a recurring validation that the stablecoin remains within the regulatory perimeter. Simultaneously, treasury downloads and archives the issuer's monthly reserve report, which must be signed by a Public Company Accounting Oversight Board (PCAOB)-registered accounting firm and confirm a liquidity ratio of at least one hundred percent in high-quality liquid assets. The combination of regulatory registry verification and attested reserve adequacy establishes the baseline: the asset is still legally and operationally what it claimed to be at the time of initial classification.

- **Real-Time "De-Peg" Monitoring.** Classification as a cash equivalent assumes that the asset's value is fixed. But secondary markets can temporarily diverge from par—a phenomenon known as de-pegging. To address this, treasury systems now incorporate automated monitoring tools that track the secondary market price of accepted stablecoins in real time. A price deviation of more than two percent—falling to $0.98 or rising to $1.02—that persists for longer than one hundred twenty minutes triggers an automatic reclassification. The asset is moved from "Cash Equivalent" to "Digital Asset (Fair Value)" until price stability is restored for forty-eight consecutive hours. This mechanical rule eliminates judgement calls: the accounting follows the market, but the monitoring is algorithmic.

- **Custodial Risk Controls.** The security of the asset itself is governed by a set of operational controls designed to prevent unauthorised movement. Initiation of stablecoin transfers requires a two-of-three multi-signature approval process. Treasury Operations serves as the initiator; the Controller or Finance function acts as the approver; Internal Audit holds the third key as an observer or emergency backup. This segregation of duties ensures that no single individual can move funds without independent authorisation. Additionally, outbound transfers are strictly limited to whitelisted counterparty addresses that have passed Travel Rule and Know-Your-Customer/Anti-Money-Laundering screening. The control environment thus extends beyond accounting classification to the custody infrastructure that holds the assets.

- **Proof of Reserve Reconciliation.** At every month-end, the Accounting Department performs a three-way match that ties together the company's internal ledger balance, the on-chain balance verified via a block explorer, and the issuer's reported total tokens in

circulation relative to its declared reserves. This reconciliation is not merely a verification of arithmetic; it is a cross-check between the company's records, the immutable ledger, and the issuer's attestations. Discrepancies trigger immediate escalation and, if unresolved, a reassessment of the asset's classification.

The Ninety-Day Exception and the GENIUS Act Guarantee

Auditors apply one additional constraint drawn from ASC 305: if a stablecoin issuer discloses that more than twenty percent of its reserves are held in commercial paper or assets with maturities exceeding ninety days, the asset cannot be classified as a cash equivalent. Instead, it must be reported as a short-term investment at fair value. This rule ensures that the "ninety-three-day" liquidity standard applies not only to the instrument's underlying assets but also to the concentration risk within those assets.

For treasurers facing lingering auditor hesitation, a final argument has emerged from the GENIUS Act itself. Section 15.11, as proposed by the OCC, legally requires Permitted Payment Stablecoin Issuers to maintain reserves on a one-to-one basis. This federal guarantee—enforceable by law—provides the "contractual right to receive cash" that GAAP has long required for cash equivalent treatment. The legal pair is no longer a matter of issuer policy; it is a matter of statute. When treasury controls are built around that statutory guarantee, the auditor's role shifts from assessing whether the asset is safe to verifying that the controls are operating as designed. That shift—from uncertainty to verifiability—is what finally makes regulated stablecoins indistinguishable from cash in the eyes of both the accountant and the law.

Real-Time Accounting and the End of the Quarter

The FASB's due process model was designed for a world in which companies reported their financial results quarterly. A standard issued in January would be implemented by the following December. Auditors would review the results annually. The entire cycle assumed that financial information was a lagging indicator of economic reality—accurate, but by definition historical.

Real-time tokenised transactions break this assumption. When every transaction is recorded on a distributed ledger in the moment it occurs, and when that ledger is publicly verifiable, the concept of a "reporting period" becomes almost arbitrary. An investor with access to a company's on-chain transaction data would know more about its current financial position than its quarterly filing will reveal three months later.

Yet the full force of that statement remains largely theoretical. Privacy is the great inhibitor. Most companies are deeply reluctant to place their entire transaction history—margins, supplier costs, and payroll—on a public ledger for competitors to inspect in real time. The "publicly verifiable" ideal therefore applies, in practice, to discrete, high-assurance contexts: the reserves backing a regulated stablecoin, the cryptographic receipts in a supply chain transparency service, or the proofs of asset custody required by auditors. The general ledger of a corporation remains private and will likely stay that way for the foreseeable future.

The FASB is not yet in a position to require real-time reporting for the broader enterprise. Its due process would take years to produce such a standard, and the technical infrastructure for continuous auditing of private ledgers is not yet universally available. But the direction of travel is clear: the accounting standards of the 2030s will be designed for continuous reporting, at least for those asset classes and transaction types where public verifiability or auditable private infrastructure can be reliably achieved. The companies building the infrastructure for real-time settlement today are building toward a world in which their accounting standards have not yet caught up – and where the boundary between what can be made public and what must remain private will define the next generation of financial reporting.

The Missing Link: From Real-Time Data to Real-Time Trust

A blockchain provides a record of what happened. An audit proves that what happened was legal, accurate, and authorised. The transition from traditional snapshot auditing to continuous auditing is therefore the missing link between real-time data and real-time trust. The technology to record transactions in milliseconds exists; the infrastructure to verify their integrity with equal speed is still being built. And the obstacles standing in the way are not merely technical—they are structural, legal, and deeply embedded in the way the accounting profession has operated for generations.

The Oracle Problem: Off-Chain Reality

The most fundamental challenge is not tracking the digital token; it is verifying that the token represents a physical asset in the real world. A blockchain can confirm a payment for ten thousand tons of steel in real time, but an auditor still needs to know whether that steel actually exists in a warehouse, whether it matches the quality specified in the contract, and whether it has been damaged or degraded. The gap between on-chain representation and off-chain reality is the oracle problem.

The emerging solution lies in the integration of Internet of Things sensors and cryptographic oracles that feed physical data—GPS coordinates, temperature readings, and weight measurements—directly onto the ledger. When a shipment of steel is weighed at the port and that weight is recorded by a tamper-evident sensor that publishes its reading to the same distributed ledger that holds the payment instruction, the two realms begin to converge. But the infrastructure for such integration remains fragmented, and the standards for what constitutes a "verified" physical input are still contested.

Data Overload Versus Materiality

Traditional auditing relies on sampling. An auditor reviews a small percentage of transactions—often five percent or less—and extrapolates conclusions about the entire population. This method assumes that anomalies are rare and that a well-designed sample will catch them.

Continuous auditing inverts that logic. It moves toward one hundred percent population testing, examining every transaction as it occurs. The immediate consequence is noise. A system that flags every micro-transaction failure, every outlier, every deviation from expected patterns will quickly overwhelm its human overseers. Alert fatigue sets in, and the risk shifts from missing a fraudulent transaction to ignoring the alerts altogether.

The solution requires sophisticated AI filters that can distinguish between a minor technical glitch and an actual fraud and between a routine deviation and a material misstatement. The technology for such filtering exists, but the profession has not yet developed agreed-upon standards for calibrating it. What constitutes a "material" alert in a continuous audit environment remains a question of judgement – and therefore a source of liability.

Legal and Liability Frameworks

An auditor's signature on a financial statement carries immense legal weight. It represents an opinion, rendered at a point in time, about whether the financial statements present fairly, in all material respects, the financial position of the entity. The temporal boundary—the date of the audit report—is a crucial limitation on liability.

Continuous auditing dissolves that boundary. If an audit is continuous, when is the auditor liable? If a fraud occurs at 2:00 PM and the system detects it at 2:05 PM, but an investor trades on the false information at 2:02 PM, who bears responsibility? The legal frameworks that govern audit liability were designed for a world of annual opinions, not real-time dashboards. There is, as yet, no Generally Accepted Continuous Auditing Standard to define the scope of the auditor's responsibility in a world where assurance is provided in milliseconds rather than months.

Fragmented Tech Stacks

Even when the legal and technical standards converge, the installed base of corporate infrastructure imposes its own constraints. Most large enterprises operate on legacy enterprise resource planning (ERP) systems—SAP, Oracle, and their counterparts—that were not designed to communicate with public blockchains. The data that lives in these systems is accurate, structured, and auditable, but it is also siloed.

The bridge between these internal private databases and the public or consortium ledgers used for trade settlement requires middleware: expensive, custom-built software that translates transactions from one format to another, ensuring that what appears on the ledger is cryptographically linked to what exists in the ERP. This layer of translation introduces both cost and risk. A misconfigured bridge can produce a ledger that is internally consistent but bears no relationship to the company's true financial position.

The Shift in Method

The contrast between a traditional audit and its continuous counterpart is stark. A traditional audit operates post-event, on an annual or quarterly cycle, using manual sampling to verify historical compliance. Its output is an auditor's opinion report, delivered weeks or months after the period end. Continuous audit, by contrast, operates simultaneously with the transaction itself, using automated 100 per cent testing to focus on risk management and prevention. Its output is a real-time dashboard or attestation—a living document rather than a static conclusion.

Key Takeaway

The technology for real-time reporting exists. Blockchains settle transactions in seconds, oracles can bring physical data onto ledgers, and AI can filter the noise of continuous testing. But the infrastructure for real-time trust—the combination of legal frameworks, liability rules, middleware standards, and professional norms—is still being built.

The 2030s will likely see the auditor evolve from a historian who examines past transactions into a system designer who audits the code and the data flows that generate those transactions. The signature on a financial statement will no longer be a stamp of approval on a static document; it will be an attestation to the integrity of the systems that produced the continuous stream of information. That shift—from verifying the record to verifying the infrastructure that creates the record—is the final frontier of the cognitive infrastructure transformation that began with digital trade documents and now reaches into the core of financial reporting itself.

The Sustainability Schism: Single vs. Double Materiality

The most significant current divergence between the FASB and the IASB concerns not digital assets but sustainability reporting. The IASB, through its sister body the International Sustainability Standards Board (ISSB), has adopted an integrated approach. IFRS S1 and S2 require companies to report on climate-related risks and opportunities using a single-materiality lens—that is, how the world affects the company's financial value. Built on the foundation of the Task Force on Climate-related Financial Disclosures (TCFD), these standards ask a focused question: what climate-related factors could reasonably affect a company's prospects, cash flows, and access to capital? The ISSB's argument is that impacts on the world eventually become financial risks; the standards therefore concentrate on the point where environmental and social phenomena translate into enterprise value.

The European Union's Corporate Sustainability Reporting Directive (CSRD), together with the accompanying European Sustainability Reporting Standards (ESRS), goes further. It requires double materiality: companies must report not only on how the world affects them (the "outside-in" lens) but also on how they affect the world—their outward impacts on the environment, human rights, and society—regardless of whether those impacts are financially material in the short term. A company whose operations cause significant environmental

damage but face no near-term financial consequences from doing so is, under the EU framework, is required to report that damage in detail. The underlying philosophy is multi-stakeholder: the report is addressed not only to investors but also to workers, communities, and civil society.

The FASB has declined to enter this territory. Its mandate, as it interprets it, is financial reporting for capital providers, not impact reporting for society. In the United States, the primary driver of climate-related disclosure rules is the Securities and Exchange Commission, which in March 2024 issued its own climate disclosure rule—a rule that subsequently faced legal challenges and was stayed in court. The FASB itself remains a purist, sticking strictly to the financial ledger and leaving the question of what should be disclosed about sustainability to the SEC and to other standard-setters.

The result is a reporting schism that any large company with EU operations must navigate. Filing a US Form 10-K under SEC rules and FASB accounting standards means omitting the kind of detailed sustainability information that is now mandatory under the CSRD. Simultaneously, filing a European sustainability report under ESRS means compiling a separate set of disclosures that often requires different data gathering, different definitions of materiality, and a different narrative about the company's role in society. The cognitive infrastructure—once unified around the common goal of investor-oriented financial reporting—is fragmenting along geopolitical lines. Large multinationals are now maintaining two distinct reporting tracks. The US track is focused on net income and the risks that affect shareholder value. The EU track is focused on net impact—the company's effects on the planet and society, even where those effects do not yet register as financial risks. The costs of this bifurcation are not merely compliance costs; they represent a deeper fragmentation of the cognitive infrastructure that once allowed capital to flow across borders on the basis of a single, comparably audited set of numbers. In its place is emerging a patchwork of regimes, each with its own materiality lens, each asking fundamentally different questions about what matters and to whom.

IFRS S1 and S2: The Architecture of Climate Reporting

To understand how the ISSB's sustainability standards operate, it helps to view IFRS S1 as the foundational rules and IFRS S2 as the climate-specific application. Together, they establish a framework that moves sustainability reporting from the margins of corporate disclosure into the core of how companies explain their prospects. While both standards are built on the same conceptual foundation—financial materiality, or how sustainability affects the company's enterprise value—they operate at different levels of granularity.

The Four Pillars of Disclosure

Both IFRS S1 and S2 are structured around the four pillars of the Task Force on Climate-related Financial Disclosures (TCFD), a design that ensures sustainability information is not relegated to a separate "side report" but integrated into the governance, strategy, and risk management processes that drive the business.

The first pillar is governance. Companies must disclose the oversight and processes their boards and management use to monitor sustainability-related risks and opportunities. This is not a statement of policy; it is a description of who is accountable and how that accountability operates.

The second pillar is strategy. Here, the focus shifts to the risks and opportunities themselves—their nature, their timeframe, and their actual and anticipated effects on the business model, supply chain, and financial planning.

The third pillar is risk management. Companies must describe the processes used to identify, assess, and prioritise sustainability-related risks and how those processes are integrated into the organisation's overall risk management framework.

The fourth pillar is metrics and targets. This is where disclosure becomes quantitative: the metrics used to measure performance, the targets set, and the progress against those targets.

IFRS S1: The General Requirements

IFRS S1 establishes the overarching framework for all sustainability-related disclosures. Its scope extends to all material sustainability risks and opportunities—not only climate, but also labour practices, water usage, ethical conduct, and any other factor that could reasonably affect the company's prospects. The standard requires a "fair presentation" of these risks, meaning that the information must be complete, neutral, and accurate, and it must enable investors to understand the connection between sustainability matters and the company's financial statements.

In terms of metrics, IFRS S1 references the industry-specific metrics developed by the Sustainability Accounting Standards Board (SASB), providing a ready-made set of indicators for sectors ranging from banking to biotechnology. The materiality lens is single-materiality: a strike that interrupts production is disclosed because it affects revenue, not because strikes are inherently undesirable social events. The standard asks what matters to enterprise value, not what matters to society.

IFRS S2: The Climate-Specific Application

Where IFRS S1 is broad, IFRS S2 is deep. It takes the same four-pillar structure and applies it specifically to climate, introducing requirements that are both more granular and more technically demanding.

The most significant of these is mandatory scenario analysis. Companies must assess the resilience of their strategy and business model to climate-related scenarios—typically a 1.5°C or 2°C warming trajectory—and disclose the assumptions, analyses, and implications of that analysis. This requirement forces companies to model the future, not merely report the past.

IFRS S2 also mandates the disclosure of Scope 1, 2, and 3 emissions. Scope 1 covers direct emissions from owned sources; Scope 2 covers indirect emissions from purchased energy; Scope 3 covers the full value chain, including suppliers and customers. The inclusion of Scope 3—often the largest component of a company's carbon footprint—is a significant expansion of the reporting burden, requiring data collection that extends beyond the legal boundaries of the reporting entity.

Finally, IFRS S2 distinguishes between physical risks (storms, floods, chronic temperature changes) and transition risks (carbon taxes, regulatory shifts, and market preferences), requiring companies to disclose both categories and their expected financial effects.

Why Single Materiality Matters

Under IFRS S1 and S2, a company does not report an impact simply because it is "bad for the world." It reports because that impact has the potential to affect the company's financial position, cash flows, or access to capital. A social issue—such as a labour strike—is material because it interrupts production. A climate issue—such as a new carbon tax—is material because it increases operating costs. The single-materiality lens focuses relentlessly on the point where sustainability meets the income statement and balance sheet.

Implementation in 2026

By 2026, the voluntary phase of IFRS S1 and S2 has ended in several major jurisdictions. The United Kingdom, Australia, Canada, and others have moved to mandatory adoption, with requirements phased in over successive reporting periods.

A "climate first" approach has been widely adopted: in the first year of mandatory reporting, many companies are required to report only on IFRS S2 (climate) and may delay the broader sustainability disclosures required by IFRS S1 (social, governance, water, etc.) by one year. This relief acknowledges the significant data burden that full sustainability reporting imposes, particularly on Scope 3 emissions and on the collection of industry-specific metrics.

For companies caught between the ISSB framework and the EU's Corporate Sustainability Reporting Directive, a new tool has emerged. The ISSB and the European Financial Reporting Advisory Group (EFRAG) recently released an interoperability map—a technical guide that allows companies to collect data once and map it to both the IFRS and CSRD reporting formats. The map does not resolve the underlying philosophical difference between single and double materiality; a company that reports under both regimes must still produce two distinct sets of disclosures, each with its own materiality threshold. But it does reduce the duplication of effort, enabling a single data-gathering infrastructure to serve two different reporting destinations. The fragmentation of cognitive infrastructure is being managed, at least operationally, by a layer of translation that did not exist when the two regimes first diverged.

PART V — THE CONVERGENCE

The Deterministic Turn: When Rules Become Code

Social Technology to Deterministic Technology

The central argument of this paper is that we are living through a category transition in the nature of rules. For the past century, the rules governing global trade—Incoterms, letters of credit, accounting standards, and internet protocols—have been social technologies: agreements maintained by human consensus, interpreted by human experts, and enforced by human institutions. They were legible, contestable, and adaptive. When they produced unjust or unworkable outcomes, humans could challenge them through the governance processes that created them.

The convergence of blockchain-based smart contracts, agentic AI, and real-time accounting is transforming these social technologies into deterministic technologies: logic embedded in code that executes automatically when specified conditions are met. When a GPS sensor on a shipping container triggers a smart contract payment without any human intermediary, the ICC's Incoterms are no longer a rulebook that humans consult — they are an algorithm that machines execute. When an AI agent negotiates a supply agreement using MCP and executes it on a programmable ledger, the IETF's protocols are not merely the pipes through which the negotiation travels — they are the medium in which it occurs.

CRITICAL OBSERVATION

"Deterministic rules are perfectly consistent and infinitely scalable. They are also, by their nature, inflexible. Every edge case that a human expert could resolve with judgement becomes a potential system failure when the rules are encoded. The 'grey area' is not a flaw in human governance — it is a feature."

The Model Context Protocol: The Universal Adapter for Agentic AI

Think of the Model Context Protocol (MCP) as the USB-C port for artificial intelligence. Before MCP, connecting an AI agent to a shipping database, a GPS provider, and a blockchain required custom, fragile code for every single connection. Each integration was hand-built, brittle, and prone to breaking whenever the underlying system changed. The result was a landscape of point-to-point connections—powerful in isolation, but incapable of scaling to the multi-agent, multi-system orchestration that modern trade finance demands.

MCP changes this architecture at its foundation. An AI agent no longer needs to know how a shipping database structures its queries, how a GPS provider formats its location data, or how a

blockchain expects transactions to be signed. Instead, the agent speaks a standardized language to an MCP server, which sits between the agent and the underlying system. The server translates the agent's request into the native format of the database, the GPS service, or the ledger, and returns the results in a structure the agent can consume without adaptation. The agent gains universal access; the systems remain unchanged; and the complexity of integration is concentrated in the server layer, where it can be managed, audited, and secured.

From Negotiation to Execution: The Three Steps

When an AI agent negotiates and executes a trade transaction using MCP, the process unfolds in three distinct steps.

The first is context retrieval—the read operation. The agent uses MCP to pull real-time data from across the transaction's information landscape: the current spot price of freight from Shanghai to Los Angeles, the specific delivery requirements encoded in the smart contract, the vessel's GPS coordinates, and the counterparty's credit rating. MCP does not dictate what data the agent retrieves; it provides a uniform interface to data that previously required separate, incompatible connections.

The second step is autonomous reasoning. Armed with this context, the agent processes the information against its programmed objectives. The reasoning may be complex: "Find the cheapest route that meets the Incoterm 'Delivered at Place' while ensuring the carrier is on the approved counterparty list and the estimated arrival time aligns with the buyer's warehouse availability." The agent weighs constraints, evaluates options, and arrives at a decision—all without human intervention.

The third step is deterministic action—the write operation. Having decided, the agent uses an MCP tool to execute. It sends a command back through the protocol: "Execute the payment because the GPS coordinate 31.2304° N has been reached" or "Release the Bill of Lading to the buyer now that the funds are confirmed." The action is not a recommendation; it is an instruction. MCP carries it to the relevant system, and the system executes it.

The Shift in Protocol: From Carriage to Constraint

There is a subtle but profound shift embedded in this architecture. Historically, a protocol like HTTPS served as a carrier: it moved a PDF of a contract from one party to another, but the contract itself—its terms, its enforceability, its execution—lived outside the protocol. The protocol was a pipe; the intelligence was in the humans reading the document.

In the world of agentic AI and MCP, the protocol itself becomes the medium of execution. The set of tools exposed by an MCP server defines the universe of possible actions. If the server does not expose a "Cancel Transaction" tool, the AI agent physically cannot cancel the transaction—not because it lacks the authority, but because the protocol provides no path to exercise that authority. The rules of the game are literally constrained by the technical capabilities of the protocol.

This inversion—from protocols that carry information to protocols that define permissible actions—is what transforms the IETF's technical standards from mere infrastructure into a form of governance. When an AI agent negotiates and executes a trade, it does so within a landscape of possibilities that has been predefined by the MCP server's tool set, the authentication mechanisms that determine who may call which tools, and the audit logs that record every call. The agent exercises autonomy, but within boundaries that are not legal or contractual in the traditional sense; they are technical, embedded, and enforced by the protocol itself. The shift from "pipes" to "medium" is therefore not merely a technological evolution; it is a change in how rules are instantiated, enforced, and experienced in the cognitive infrastructure of global trade.

Two Worlds of Integration: Traditional APIs vs. MCP-Based Agentic Workflow

Imagine two different versions of a global supply chain transaction.

In the traditional API world, the process resembles a series of phone calls between people who speak different languages. Each system—the warehouse, the shipping company, and the bank—exposes its own custom-built adapter, known as an API. These adapters are unique to each system, shaped by their internal logic, security models, and data structures. When the shipping company updates its software, the connection breaks, and a human engineer must intervene to restore the translation. The AI in this world is a sophisticated reader: it can see the data that flows through these fragile pipes, but it struggles to act across them because every button it needs to press is shaped differently, and there is no universal hand.

The MCP-based agentic workflow represents a fundamental shift in the environment itself. Instead of a tangle of custom adapters, every system plugs into a universal power strip. The Model Context Protocol (MCP) acts as the common interface, abstracting away the idiosyncrasies of each underlying application.

When an AI agent enters this environment, it does not need a manual for the warehouse or a specialised login for the bank. It queries the MCP server, which instantly provides a menu of available tools—CheckInventory, CalculateTax, AuthorizePayment—each presented in a uniform language that the agent understands natively. The agent no longer merely observes data from disparate sources; it inhabits a contextual bubble where the data and the tools to manipulate that data are fused into a single, coherent interface.

This architecture gives rise to a new narrative of autonomy. In the traditional model, the AI is a passenger, reporting what it sees to a human driver who must manually coordinate across systems. In the MCP model, the AI is the driver. It identifies a delay in a Suez Canal shipment through a real-time data tool, reasons that a contractual penalty is imminent, and—using the same protocol—automatically negotiates a secondary shipment and signs the digital ledger to fund it. The entire loop—perception, reasoning, action—occurs within a unified protocol environment.

The pipe has not simply become faster; it has transformed into an interactive dashboard where the AI's thoughts and its actions happen in the same breath. The transition from social to deterministic—first observed in the shift from human-consensus rules to machine-executable

logic—finds its final expression here. The human is no longer the bridge between systems; the protocol is the bridge, and the AI is the traveller crossing it. What remains of human involvement is not in the loop but at the edges: setting the objectives, defining the constraints, and auditing the results. The execution itself now belongs to the infrastructure.

The Deterministic Incoterm: CIF in the Age of Agentic AI

To illustrate the shift from social to deterministic logic, consider a classic Incoterm: CIF—Cost, Insurance, and Freight. In the traditional world, if a ship sinks, a small army of claims adjusters, lawyers, and bank agents spends months arguing over the precise moment risk transferred. The outcome depends on human interpretation, memory, and negotiation. In the deterministic world of agentic AI and the Model Context Protocol (MCP), that social argument is replaced by logical execution—a transaction that resolves itself in seconds because the rules are embedded in the infrastructure, not merely written in a rulebook.

The Scenario: A Mid-Ocean Loss

Imagine a cargo of high-end semiconductors travelling from Taiwan to Rotterdam. The buyer's procurement officer is not a human but an AI agent, tasked with monitoring the shipment in real time. It maintains a constant MCP connection to a satellite-linked IoT sensor embedded in the container stack on the vessel. The agent does not wait for reports; it receives data continuously.

While the ship traverses the Indian Ocean, a severe storm causes a container stack to collapse. The IoT sensor on the semiconductor crate detects two events in rapid succession: a high-impact jolt, followed by saltwater immersion. Within milliseconds, the sensor publishes these readings to the MCP server, which makes them available to any authorised agent.

Agentic Reasoning: The Medium as the Arbiter

The buyer's AI agent, through its MCP server, instantly pulls the telemetry data. It does not merely forward the information to a human for review; it evaluates the data against the terms of the transaction, which have been encoded as a smart Incoterm on a programmable ledger. The underlying logic is drawn from the CIF rule: the seller's risk ends once the goods are delivered on board the vessel. The agent knows that the impact and immersion occurred after that point, meaning the risk has already passed to the buyer.

This reasoning is not a matter of interpretation. The terms are expressed in deterministic code, and the telemetry is cryptographic proof. The agent does not weigh equities or consider goodwill; it executes the transaction as it is programmed.

Deterministic Action: Execution Without Friction

In a social-technology world, the buyer might attempt to stall payment or demand a discount, arguing that the goods are ruined. A dispute would follow, consuming weeks or months. In the deterministic world, the AI agent recognises that the conditions for payment have been met. It does not choose to pay; the protocol requires it.

The agent executes a "Release Funds" command on the blockchain ledger because the telemetry proves the seller fulfilled their obligation. Simultaneously, the agent uses other MCP tools to open a claim with the digital insurance provider—since CIF requires the seller to procure insurance for the buyer—and to search for a replacement spot-market shipment of semiconductors to prevent a factory shutdown. The entire sequence occurs without human intervention.

The Result: A Binary Switch, Not a Dispute

The dispute never actually happens. There is no courtroom, no letter of credit discrepancy, no debate over whether the packaging was adequate or the storm was foreseeable. The Incoterm has shifted from a set of guidelines for humans to argue over into a binary switch—a zero or one—in a machine's decision-making loop. The AI did not simply use the protocol to talk; it operated inside the protocol's constraints to settle a multimillion-dollar event in seconds. The social process that once defined trade has been replaced by a deterministic one, and the cognitive infrastructure that governs the transaction is no longer a rulebook to be interpreted but code to be executed.

The Era of Executable Law

The shift from social to deterministic technologies represents the final decoupling of global commerce from local jurisdiction. When rules are maintained by human consensus, they are inevitably subject to the friction of culture, language, and legal interpretation. A contract clause that reads "reasonable dispatch" means one thing in London and another in Shanghai; a letter of credit that requires "clean on board" notation leaves room for a dozen documentary disputes. This friction, for all its inefficiency, has historically provided a human buffer—a space for mercy, for context, for the kind of common sense that no rulebook can fully capture.

Moving those rules into the medium of AI agents and programmable ledgers trades that buffer for total certainty. The buffer is replaced by a boundary: the transaction executes exactly as coded, without interpretation, without delay, and without appeal to a judge's sense of fairness. The result is a form of law that is not merely digitised but executable—embedded in the infrastructure of commerce rather than stored in the libraries of courts.

The New Infrastructure of Trust

We are moving toward a global operating system whose properties are fundamentally different from the one that preceded it.

Compliance is no longer a choice. In the paper-based world, a party could choose to comply, to negotiate an exception, or to litigate. In the deterministic world, compliance is a mathematical requirement of the protocol. If the terms are not met, the transaction does not proceed—not because a bank officer rejects the documents, but because the code has no path to execute.

The expert is being replaced by the algorithm. The lawyer who drafts the clause, the accountant who verifies the entry, the trade finance specialist who reconciles the documents—these roles are increasingly performed by MCP-enabled AI agents that read code as fluidly as a human reads a letter. The agent does not interpret; it executes. Its expertise lies not in judgement but in the precision of its logic.

Transparency becomes real-time. The books are no longer a static document, updated quarterly, offering a backward-looking portrait of the enterprise. They become a live, streaming reflection of every physical move a shipping container makes, every transfer of title, and every release of payment. The balance sheet and the supply chain converge into a single, continuously verifiable flow of information.

The Final Category Transition

The paper's central argument concludes that we are not merely digitising old paper processes. We are changing the ontology of a rule. A rule is no longer an ought—" You ought to pay when the goods arrive"—but it is the payment that is triggered when the GPS coordinate is reached. The distinction between obligation and execution collapses. What was once a promise becomes a condition; what was once a judgement becomes a computation.

As agentic AI becomes the primary actor in this landscape, the human role shifts from operator to architect. We will no longer spend our time enforcing rules in the thick of the transaction. Instead, we will spend our time auditing the code that enforces them for us—designing the constraints, verifying the logic, and stepping in only when the system signals a frontier that its original parameters cannot resolve. The social technology of the last century served global commerce well, providing a framework of trust built on reputation, relationship, and the slow accumulation of precedent. But the speed of the next century requires a deterministic foundation that only code, blockchain, and AI can provide. The infrastructure is being rewritten now, and the practitioners who understand that transformation will be the ones who shape the rules of the world that follows.

The Black Box Problem

The ICC, IETF, and FASB/IASB are all, in the paper's original framing, "glass box" systems. Their committee notes are public. Their reasoning is documented. When the FASB issues a standard requiring fair value accounting for digital assets, it publishes a Basis for Conclusions document explaining every significant decision. When the IETF publishes an RFC, the mailing list debates that produced it are permanently archived. When the ICC updates Incoterms, the national committee consultations that shaped the revision are, at least in principle, discoverable. The process is as transparent as the outputs.

Algorithmic governance is structurally different. The "black box" problem is often an issue of semantic opacity rather than literal secrecy. A smart contract's code may be open-source and visible on a ledger, but it lacks the Basis for Conclusions that anchors traditional standards. We can see the execution logic, yet the legal intent—the trade-offs made during programming, the interpretive choices embedded in the code—remains invisible. When the code becomes the sole arbiter of the rule, we lose the ability to verify whether the machine's execution aligns with the institution's original authoritative intent.

This opacity deepens with AI agents. An agent trained to negotiate using ICC Incoterms has absorbed a model of those terms from its training data, but whether that model accurately reflects the authoritative interpretation of the ICC Banking Commission is not guaranteed. An algorithmic audit engine applying FASB standards to a real-time transaction ledger is applying a model of those standards, not the standards themselves. The mapping from the authoritative text to the executable logic is neither transparent nor verifiable without deep expertise in both the subject matter and the machine learning architecture that produced the mapping.

The risk is not that these algorithmic systems will produce obviously wrong results. It is that they will produce subtly wrong results—results that look correct, are executed at scale, and create systemic distortions that no individual transaction would reveal. The capacity for meaningful oversight disappears not because the code is hidden, but because the reasoning that would allow us to judge its correctness has been replaced by execution alone. The glass box of the standard-setter becomes a ledger whose contents are visible, but whose logic—disconnected from the documented intent of the rule-makers—is, for most observers, effectively opaque.

The Semantic Gap: When Code Replaces Reasoning

In the "glass box" systems of the ICC, IETF, and FASB, ambiguity is managed through documentation. When a rule is contested, the participants return to the archives—the Basis for Conclusions, the mailing list threads, and the national committee minutes—to recover the intent behind the words. The rule exists in a space where interpretation is expected and where the reasoning that produced the rule remains accessible.

Algorithmic governance collapses that space. The ambiguity does not disappear; it is solved at the moment the code is compiled, often without the developer realising they have made a legal choice. A smart contract's logic may be visible on a ledger, but the normative metadata—the "why" that anchors traditional standards—is absent. We can observe how the execution

proceeds; we cannot recover why the programmer set a particular threshold, omitted a fallback, or encoded one interpretation of an Incoterm over another.

This is the semantic gap: the distance between what a programmer thinks a rule means and what a legal body actually intended. The developer becomes, in effect, a legislator without a licence, translating complex ICC or FASB guidance into binary logic under commercial pressure, with no published Basis for Conclusions to accompany their work.

The risk is not that a single contract misfires. It is that a million contracts, all executing a slightly flawed interpretation of “fair value” or “reasonable dispatch", create systemic drift before any human auditor notices. The rule, stripped of its documented intent, becomes indistinguishable from the code. And when the code is the sole arbiter, meaningful oversight vanishes—not because the code is hidden, but because the reasoning that would allow us to judge its correctness has been replaced by execution alone.

The Semantic Gap in Practice: When CIF Meets Code

In a traditional glass-box environment, the ICC provides extensive guidance on when delivery occurs and when risk passes from seller to buyer. The rules are not merely text; they are accompanied by Guidance Notes, Opinions, and decades of documented practice. When a dispute arises, a judge does not interpret the Incoterm in isolation. They consult the archives—the institutional memory of how the rule was intended to operate. The reasoning is recoverable. The intent is discoverable.

Now consider a smart contract designed to automate a CIF (Cost, Insurance, and Freight) transaction. A developer, working under commercial pressure and without the benefit of a Basis for Conclusions, writes the code. The requirement seems straightforward: trigger payment upon "delivery".

Under the ICC’s authoritative guidance, delivery in a CIF contract occurs when the seller places the goods on board the vessel—not when the vessel reaches the destination. Risk passes to the buyer at that same moment. The seller’s obligation is to ship; the buyer’s obligation is to pay, regardless of what happens during transit. If the ship sinks mid-ocean, the buyer must still pay and then claim against the insurance that the seller was required to procure.

The developer, however, interprets “delivery” differently. They hook the smart contract into a GPS API, reasoning that delivery should be tied to the physical arrival of the goods at the destination port. Payment triggers when the ship reaches Rotterdam, not when it leaves Kaohsiung. The code compiles. The contract deploys.

The Subtly Wrong Execution

A ship carrying high-value semiconductors departs Taiwan for Rotterdam. Mid-voyage, a storm causes a container stack to collapse. The vessel never reaches the destination port. The GPS API reports that the coordinates have not been attained. The smart contract, following its logic, determines that delivery has not occurred. It automatically refunds the buyer and releases the collateral.

The code is not broken. It executed exactly as the programmer instructed. But it violated the legal definition of the CIF term. The seller, having placed the goods on board in good order and procured the required insurance, now finds the payment reversed. The insurance claim is technically available, but the buyer has no incentive to pursue it; the buyer has been made whole by the refund. The seller must now initiate legal proceedings to recover the funds—proceedings that may take years and that must overcome the "authoritative" result produced by the algorithm.

Why This Is a Black-Box Problem

A third-party observer examining the ledger sees a successful execution: a refund, cleanly processed, with a clear audit trail of the GPS coordinates. They do not see that the execution violated the legal definition of the term. The logic is opaque not because the code is hidden but because the reasoning that would reveal the error—the documented intent of the ICC, the distinction between shipment and arrival, and the allocation of risk in CIF—is not present in the code or its outputs.

Recourse becomes a paradox. To challenge the algorithm, the seller must exit the deterministic environment and enter a traditional court, arguing that the code misapplied the rule. But the rule, in this context, is the code. The court must interpret not only the Incoterm but also the relationship between a legal standard and its algorithmic implementation—a task for which no established framework yet exists.

The deeper risk is systemic. If this GPS-arrival logic is copy-pasted into thousands of maritime smart contracts, the industry quietly drifts away from the ICC's authoritative standards. A new standard emerges—call it the Developer Standard—shaped by the assumptions of programmers rather than the deliberations of the ICC Banking Commission. This shift occurs without a single public debate, without a committee note, without any of the institutional guardrails that have governed global trade for a century. The glass box of the standard-setter is replaced by the opaque logic of the compiler, and the semantic gap becomes a structural feature of the new cognitive infrastructure.

The Case of the Sinking Ship: A Tale of Two Standards

In a traditional ICC-governed contract, the CIF Incoterm specifies that the seller's delivery is complete—and risk transfers to the buyer—the moment the goods are loaded on board the vessel at the port of shipment. The logic is drawn from decades of documented practice: Guidance Notes, national committee consultations, and the institutional memory of the ICC Banking

Commission. The trigger event is physical and unambiguous: the goods cross the ship's rail at the port of origin.

Now consider a smart contract designed to automate the same CIF transaction. The developer, working without the benefit of a Basis for Conclusions, interprets "delivery" through the lens of a GPS API. The contract is written to trigger payment when the ship's coordinates match the destination port. The code is clean, the API is reliable, and the contract deploys without incident.

The difference between the two models is not a matter of broken code versus functional code. It is a matter of interpretation. Under the ICC standard, the trigger event is the loading at origin. Under the smart contract model, the trigger event is arrival at destination. The gap between them is the semantic gap—the distance between what the rule says and what the programmer assumes.

The Subtly Wrong Execution

A vessel carrying high-value cargo departs from Kaohsiung bound for Rotterdam. Mid-voyage, a storm causes the ship to sink. The GPS coordinates never reach the destination port. The smart contract, following its logic, determines that delivery has not occurred and automatically refunds the buyer.

Under the ICC standard, the buyer would still be liable to pay. The risk transferred at the port of origin; the buyer now holds the insurance policy and must make a claim against it. The smart contract, however, has executed a refund. The seller—who loaded the goods in good order, procured the required insurance, and fulfilled every obligation—now finds the payment reversed. The code is not broken. It did exactly what the programmer told it to do. But what the programmer told it to do did not reflect the rule.

Why This Is Structurally Different

The failure here is not technical; it is interpretive. The developer equated "delivery" with "arrival", an equation that the ICC's authoritative guidance explicitly rejects. In a glass-box system, this error would have been caught during the drafting phase. A lawyer would have flagged the discrepancy by consulting the Basis for Conclusions; a trade finance specialist would have noted that CIF risk transfers at the ship's rail, not at the destination.

In a smart contract, no such review occurs. The "rule" is simply the if/then statement written by the developer. There is no committee note, no public debate, no archival record of the interpretive choice. The semantic drift becomes embedded in the code and, once deployed, becomes the effective rule for that transaction.

The consequences extend beyond the individual contract. If this specific smart contract template—the GPS-arrival logic—is copy-pasted across thousands of maritime transactions, the industry begins to shift, silently and without deliberation, toward a destination-based risk model. The ICC's carefully negotiated allocation of risk, designed to encourage the development

of robust insurance markets, is quietly replaced by a developer's assumption about what "delivery" means. This shift occurs without a single public debate, without a committee note, without any of the institutional guardrails that have governed global trade for a century.

Loss of Context

The smart contract—and the AI agents that interact with it—lack institutional memory. They do not know why the ICC chose to place risk on the buyer during transit. They do not know that the rule was designed to align incentives, to ensure that the party with the most control over the goods during the voyage (the buyer) also bears the risk and therefore purchases insurance. They only know the binary state of a GPS ping. The context that gives the rule its meaning is absent from the execution.

Toward a Basis for Code

For algorithmic governance to be legitimate—to serve the same function as the glass-box systems that preceded it—it must adopt the paper-trail DNA of its predecessors. Open-source code is necessary but not sufficient. We need more than visibility into the logic; we need traceability from the legal standard to its implementation.

What is required is a Basis for Code: a documented link between the authoritative rule—the ICC's Guidance Note, the FASB's Basis for Conclusions, and the IETF's mailing list debate—and the logic gates that instantiate it. This documentation must be as accessible, as auditable, and as subject to public scrutiny as the committee minutes that produced the underlying standards. Without it, the shift from social to deterministic technologies risks not only the loss of efficiency but the loss of accountability itself. The code will execute with perfect fidelity to its programming, but we will have no way of knowing whether what it executes bears any relationship to the rules we thought we were writing.

The transition from human-interpreted standards to machine-executable logic does not have to result in a loss of oversight. However, it requires a new archival requirement: the Technical Basis for Conclusions (TBC). Just as the FASB or IETF provide a narrative record of their debates and decisions, algorithmic systems must be accompanied by a formal mapping that links every 'if/then' statement to its authoritative source in the ICC or GAAP frameworks. By requiring that code be 'annotated' with its legal and economic intent, we can transform the algorithmic black box back into a glass box—ensuring that when the rule is the code, the reasoning remains discoverable, contestable, and human-centric.

The Case for Hybrid Governance

The appropriate response to the deterministic turn—a phrase that captures the philosophical shift from rules maintained by human consensus to rules embedded in executable code—is not a return to paper-based trade finance. The efficiency gains from digital standards are real and substantial: reduced transaction costs, compressed settlement cycles, and democratised access to trade finance. A world in which Terra Verde, an Ethiopian coffee cooperative, can access

revolving credit at institutional rates on the basis of a verified digital trade footprint is unambiguously better than a world in which it cannot.

The appropriate response is hybrid governance: automated execution with human exception-handling, algorithmic efficiency with institutional oversight, and deterministic rules with contestable governance. The ICC's Hybrid Bridge Clause offers a template for this approach at the level of individual transactions. What is needed now is a set of institutional arrangements that apply the same logic at the systemic level.

Concretely, this means maintaining the human governance bodies—the ICC's commissions, the IETF's working groups, and the FASB's board—not as slow alternatives to automated systems but as the constitutional authorities that define the parameters within which automated systems operate. The machine executes the rule. The human institution defines what the rule means, identifies when it is producing perverse outcomes, and updates it through a process that is legitimate, transparent, and contestable.

The deterministic turn does not render these institutions obsolete. It elevates their role. In a world where code executes without human intervention at every step, the quality of the governance that sits above the code—the standard setting, the interpretive guidance, the mechanism for recourse—determines whether the system serves its users or simply executes its programmers' assumptions. The institutions that built the cognitive infrastructure of the twentieth century are now charged with architecting the governance layer of the twenty-first. Their legitimacy, transparency, and willingness to adapt will decide whether the handshake between machines remains tethered to the rule of law—or whether the law becomes whatever the code says it is.

PART VI — IMPLICATIONS FOR DIGITAL TRADE FINANCE

The $2.5 Trillion Problem and Its Structural Solution

The Trade Finance Gap Reconsidered

The $2.5 trillion global trade finance gap is not a market failure in the conventional sense. It is not the result of insufficient capital—global bank balance sheets contain ample liquidity. It is the result of information asymmetry: banks cannot efficiently evaluate the creditworthiness of small and medium-sized enterprises in emerging markets, so they price the uncertainty into prohibitive risk premiums or decline to lend at all.

The digital standards discussed in this paper address this information asymmetry at its root. A company with three years of verified, tamper-proof shipment data recorded on a platform compliant with the ICC's URDTT (Uniform Rules for Digital Trade Transactions) has a demonstrable track record that a traditional bank relationship can never provide. An Ethiopian coffee cooperative that has completed two hundred shipments without a single dispute, documented on a distributed ledger, is a better credit risk than its paper-based balance sheet suggests. But the data alone is not enough. For that digital footprint to become legally legible to a bank in London, it must rest on a foundation of law.

That foundation is now in place. The UNCITRAL Model Law on Electronic Transferable Records (MLETR) has been adopted by a growing number of jurisdictions—the United Kingdom, Singapore, the United Arab Emirates, and others—establishing that a digital record can carry the same legal weight as a paper document. When the cooperative's shipment data is recorded on an MLETR-compliant platform, the bank can treat the digital record as a legally binding transferable record. The information asymmetry collapses not because the data is visible, but because it is enforceable. What was once a "nice to have" audit trail becomes a legally recognised track record that a bank can rely on for credit decisions. The digital infrastructure, for the first time, makes the creditworthiness of small-scale exporters legible to global finance—and the legal infrastructure ensures that legibility is backed by the rule of law.

Metric	Paper Trade (Pre-2024)	Digital / Hybrid Trade (2026)
SME Approval Time	20+ days	24–48 hours is possible for "flash" financing, but for large international trade lines, KYC/AML regulations still often push this to 3–5 days even in 2026
Transaction Costs	~3.0% of trade value	~0.7% of trade value
Document Processing	25 days average	Less than 24 hours
Insurance Premium	1.0–1.5% of cargo value	0.6–0.8% of cargo value
Working Capital Cycle	30–90 days	Near real-time, in the tech's capability, while the physical movement of goods and customs clearance still creates a floor for how fast capital can truly recycle.
SME Market Access	Local / regional only	Global, direct-to-consumer

Tokenised Invoices and the Instant Liquidity Revolution

The most immediately transformative application of digital standards to the trade finance gap is the tokenisation of invoices and receivables. In traditional trade finance, an SME waiting sixty or ninety days for an invoice to be paid faces a liquidity gap that forces it to either forgo growth opportunities or borrow at rates that reflect the bank's uncertainty about the invoice's value. The invoice is an asset—the buyer has acknowledged the debt—but it is an illiquid asset, trapped in a bilateral relationship.

Tokenised invoices convert this illiquid asset into a marketable instrument. When an invoice is issued as a digital asset conforming to internationally recognised data standards—such as ISO 20022 for financial messaging or the UN/CEFACT standards for cross-border trade—and recorded on a distributed ledger, it can be sold on secondary markets to investors who can verify its legitimacy and track record in real time. The SME receives immediate liquidity. The investor receives a short-duration asset backed by a verifiable commercial obligation. The bank provides the governance framework that makes the transaction legally sound.

The ICC's legal infrastructure is essential to making this work. A tokenised invoice is only as valuable as the legal enforceability of the obligation it represents. If that obligation is governed by recognised ICC rules—specifically the Uniform Rules for Digital Trade Transactions

(URDTT)—it provides a high-level framework for both the financial and commercial components of the trade. While the UCP 600 remains the gold standard for Letters of Credit, the URDTT is designed for the 'open account' environment where tokenised invoices thrive. When these rules are applied within an MLETR-aligned jurisdiction, the digital record is recognized as a 'transferable document' by law. This synergy standardises the risk profile of the financial instrument, allowing investors to bypass bespoke legal analysis and treat the token as a high-quality, liquid asset.

Underpinning everything is the UNCITRAL Model Law on Electronic Transferable Records (MLETR). Without MLETR, a token is merely a piece of data—a convenient representation of a receivable but not its legal equivalent. With MLETR, that token becomes the functional equivalent of a paper promissory note or a bill of exchange, recognised by law as a transferable record that can be possessed, endorsed, and enforced across borders. The standardisation of the underlying trade documentation—through URDTT and MLETR—standardises the risk profile of the financial instrument built on top of it, turning illiquid invoices into assets that global capital markets can confidently trade.

Deep-Tier Supply Chain Finance

Perhaps the most structurally significant implication of digital trade standards is their potential to extend supply chain finance beyond the first tier. In conventional supply chain finance, large buyers provide their direct suppliers with access to financing based on the buyer's credit rating. A major retailer's approved payables programme allows its tier-one suppliers to access cheap capital. But those tier-one suppliers' own suppliers—the tier-two and tier-three participants who are often the smallest and most financially vulnerable actors in the chain—receive no benefit.

The ICC's "deep-tier" supply chain finance models address this by allowing the credit quality of the large buyer to cascade through the supply chain via programmable digital instruments. When a Japanese automotive manufacturer issues a Digital Promissory Note governed by ICC rules, that note can be subdivided and passed down through the supply chain, with each tier's suppliers using their portion as collateral for local financing. Because the note is anchored in a globally recognised legal framework, a local bank in São Paulo or Nairobi can evaluate its legal quality without needing a relationship with the Japanese manufacturer.

The numbers are significant: if deep-tier supply chain finance were deployed at scale across global manufacturing supply chains, the addressable market exceeds the current $2.5 trillion trade finance gap. The constraint is not capital but legal infrastructure—the standardisation of digital instruments and the legal recognition of electronically transferable records across jurisdictions. This is precisely where the ICC's patient, decade-long governance work is paying dividends, though the dividends are still being realised. Adoption of the UNCITRAL Model Law on Electronic Transferable Records (MLETR) has accelerated—the UK's Electronic Trade Documents Act 2023 was a milestone—but many jurisdictions are still in the process of legislating the framework into local law. Until legal recognition is truly global, the deep-tier model operates at the pace of national legislative cycles.

Equally important is the challenge of interoperability. The programmable digital instruments that carry credit down the supply chain must function across different blockchain platforms, distributed ledgers, and bank-operated infrastructure. The ICC's Digital Standards Initiative (DSI) is working to ensure that these "walled gardens" can speak to one another so that a Digital Promissory Note issued on one platform can be recognised, subdivided, and financed on another. Without that interoperability, the deep-tier vision remains fragmented by technical silos. With it, the cascading credit model can scale to the full depth of global supply chains, unlocking liquidity for the smallest participants—the tier-three fabric supplier in Bangladesh or the tier-two fastener manufacturer in Mexico—who have historically been invisible to formal trade finance.

PART VII — DESIGN LESSONS

Building the Infrastructures of Tomorrow

Resilience Is Not a Formula

The comparative analysis of the ICC, IETF, and FASB/IASB reveals that resilience in cognitive infrastructure is not a single strategy but a family of strategies, each adapted to a specific combination of technical environment, stakeholder structure, and risk profile.

The ICC is most resilient against market panic and fleeting trends. Its deep institutional inertia—the result of a century of practitioner-driven consensus—makes it almost immune to regulatory fashion. When blockchain startups in 2017 declared that smart contracts would make the ICC irrelevant within five years, the ICC did not panic. It began the quiet, meticulous work of writing legal rules that would govern smart contracts. By 2026, those rules exist, and the smart contract platforms that ignored the ICC's legal infrastructure have discovered that their transactions are unenforceable in national courts. Yet this institutional inertia has not prevented the ICC from being surprisingly proactive where it matters most: environmental, social, and governance (ESG) standards. It has moved with unusual speed to integrate green trade finance rules into its latest frameworks, recognising that sustainability is not a passing fashion but a structural shift in global commerce.

The IETF is most resilient against centralised control and technological stagnation. Its radical openness—anyone with an email address can join a working group—means that new ideas cannot be suppressed by incumbents. Its insistence on "running code" means that theoretical arguments without practical implementation cannot capture the standards process. The IETF's vulnerability, as the "New IP" episode demonstrated, is to well-resourced actors who attempt to introduce proposals that serve narrow interests under the guise of technical improvement. The rough consensus model held in that case, but it required vigilance.

The FASB/IASB is most resilient against casual change and regulatory arbitrage. Its legal embeddedness—the SEC's endorsement of FASB standards and the EU's adoption of IFRS—means that companies cannot simply ignore accounting rules without regulatory consequence. Its vulnerability is political: when accounting standards threaten powerful economic interests, those interests seek congressional or regulatory intervention.

The Convergence of Three Pillars: RWA Tokenization in 2026

The most consequential test of the three institutions' resilience lies in the tokenisation of real-world assets (RWA)—the process of converting physical goods, invoices, or property into digital tokens that can be traded, financed, and settled on distributed ledgers. Each pillar brings a distinct form of infrastructure to the table, and the friction between them reveals the contours of the emerging cognitive infrastructure.

The ICC vs. IETF: The Legality-vs-Logic Bridge

The fundamental challenge in RWA tokenisation is ensuring that the digital token remains legally tethered to the physical object it represents. A tokenised shipping container, for example, must be more than a string of code; it must be recognized by a port authority, a customs official, and a bank as the functional equivalent of a paper bill of lading.

This is where the ICC and the IETF meet—and often clash. The ICC's Digital Standards Initiative (DSI) has been working with technical bodies to align IETF-style protocols—for identity, secure transmission, and data integrity—with the legal framework of the UNCITRAL Model Law on Electronic Transferable Records (MLETR). The goal is to create a stack where the protocol handles the movement of the token and the rulebook handles the consequences if the physical asset is damaged, stolen, or disputed.

The friction is cultural. The IETF moves fast, favours permissionless experimentation, and measures success in "running code". The ICC operates at the speed of practitioner consensus, insists on detailed rulebooks, and measures success in legal certainty. By 2026, a hybrid layer has emerged: IETF protocols handle the secure, atomic transfer of the token, while ICC rules provide the dispute resolution and commercial context that make the transfer enforceable across jurisdictions. The two are no longer competing; they are being forced into complementarity by the market's demand for both speed and legal safety.

FASB / IASB: The Gatekeepers of the Balance Sheet

Even when a token is technically secure (IETF) and legally valid (ICC), a company will not touch it unless the accounting standard setters say how to record it. The FASB and IASB therefore act as the gatekeepers of corporate adoption. For years, their silence on digital assets forced companies to treat tokenised receivables as intangible assets—a classification that made them illiquid and operationally awkward.

The shift in 2026 is twofold. First, the FASB has added projects to clarify derecognition—the rules under which a company can officially remove a tokenised asset from its balance sheet after it has been sold. This is critical for securitization: if the asset stays on the books, the transaction is a secured loan, not a true sale. Second, the accounting boards are under intense pressure from financial institutions to classify certain stablecoins and regulated RWA tokens as cash equivalents rather than intangible assets. The stakes are enormous: if the FASB yields, tokenized assets become treasury-grade instruments, triggering a wave of corporate adoption. If it resists, tokenization remains a niche financial product, confined to the margins of the balance sheet.

The conflict is not technical but political. The accounting boards are designed to resist lobbying, but the political pressure to treat digital dollars as dollars—and tokenised collateral as cash—has never been higher. The outcome will determine whether RWA tokenisation scales from pilot projects to the core of corporate treasury management.

The Resilience Stress Test

As these three forces converge on RWA tokenisation, their distinctive "families of strategy" become visible.

The IETF focuses on standardising the technical plumbing—protocols like ERC-7518 for on-chain identity, secure asset transfer, and interoperability between ledgers. Its resilience play is structural: by keeping the infrastructure open and permissionless, it prevents any single actor, even a dominant bank or platform, from owning the "Internet of Value".

The ICC focuses on standardising the commercial rules that make digital trade documents legally effective. Its resilience play is jurisdictional: by anchoring its rules in MLETR and ensuring that a tokenised bill of lading is recognized by port authorities in Singapore and Rotterdam, it creates a global legal layer that transcends any one national court.

The FASB and IASB focus on standardising the audit trail and tax treatment. Their resilience play is accountability: by imposing rigorous rules for when assets can be recognized, derecognised, or revalued, they prevent companies from using tokenisation to hide debt or inflate asset values—the "Anti-Enron" shield that makes the system trustworthy.

The 2026 Conflict: Speed vs. Certainty

The tension running through all three is the perennial conflict between speed and certainty. The IETF wants interoperability now—protocols that allow any agent to transact with any other, regardless of jurisdiction or legal nuance. The ICC wants to ensure that no one gets sued—that the rules are clear enough to protect banks, carriers, and insurers. The FASB wants to ensure that the numbers don't lie—that financial statements reflect economic reality, not the creative structuring of tokenised assets.

The winner of this era is not any one organisation. It is the emerging class of interoperability frameworks that successfully satisfy all three demands: technical security that meets IETF standards, legal enforceability anchored in ICC rules and MLETR, and accounting transparency that aligns with FASB and IASB requirements. These frameworks are still being built, but their architecture is now visible. They treat the token not as a substitute for law, but as a medium through which law, code, and accounting can be simultaneously satisfied. The handshake between machines, in this vision, does not bypass human institutions—it encodes their logic in a form that machines can execute and humans can trust.

The Design Principles That Transcend the Specimens

Across all three cases, a set of design principles emerges that appears necessary—though not sufficient—for cognitive infrastructure to remain resilient across technological and political disruptions. From a functionalist perspective, these principles ensure durability; however, a critical lens reveals that this very resilience can be a product of exclusion, raising the question of for whom these infrastructures are designed to endure.

Align Incentives Before Writing Rules

The UCP 600 is adopted by banks worldwide, not because the ICC has the power to compel adoption, but because banks that do not adopt it face costly, unpredictable litigation. The IETF's protocols are implemented by router manufacturers not because anyone mandates it but because devices that do not implement them cannot communicate with the rest of the internet. The first question for any cognitive infrastructure designer is not "what rules do we need?" but "why would anyone follow these rules, and how do we ensure that incentive persists even when the rules are inconvenient?" Yet this principle, focused on the self-interest of existing actors, risks embedding the priorities of incumbents. An incentive structure that compels adoption by global financial institutions does not guarantee that the resulting rules serve the needs of emerging economies or smaller stakeholders.

Build Conflict Resolution Before You Need It

The ICC's International Court of Arbitration was not an afterthought. It was the product of a recognition that rules without enforcement are suggestions. The IETF's "rough consensus" process is not merely a decision-making tool—it is a conflict resolution mechanism that prevents unresolved technical disagreements from fragmenting the standards process. The FASB's Basis for Conclusions documents create a public record of the reasoning behind standards that constrains future interpretations and provides grounds for challenge. However, conflict resolution mechanisms are only as equitable as the access to them. A court of arbitration or a due process that requires significant legal or technical resources can resolve disputes effectively among well-resourced participants while remaining functionally inaccessible to those without such means, effectively resolving conflicts by silencing dissent.

Design for Change, Not for Permanence

The most durable cognitive infrastructures are those with the most sophisticated mechanisms for change, not those that resist change most fiercely. The IETF's RFC process allows any standard to be obsoleted by a successor RFC. The ICC publishes Incoterms on a decade-long cycle that is slow enough to ensure stability but regular enough to prevent calcification. The FASB's due process, though slow, provides a legitimate pathway for any stakeholder to challenge any standard. Institutions that resist change accumulate contradictions until they become brittle. Institutions that manage change through transparent, legitimate processes accumulate credibility. But a legitimate pathway is not the same as an equitable one. When change mechanisms require sustained engagement from sophisticated, well-funded actors, they favor

incremental adjustments that preserve the status quo over the more fundamental overhauls that developing nations or marginalised groups might require.

Preserve the Human Layer—And Broaden It

The most consequential lesson of the digital transition is the one most easily overlooked: automated execution requires human oversight not despite its efficiency but because of it. A smart contract that executes perfectly but on the basis of incorrect inputs causes harm at a scale and speed that no human institution can prevent after the fact. The human governance layer—the ICC's commissions, the IETF's working groups, and the FASB's board—is not an obstacle to digital efficiency. It is the constitutional authority that makes digital efficiency legitimate and correctable.

Yet this human layer is not without its own vulnerabilities. The functionalist view treats this layer as a neutral source of stability, but a critical perspective recognises it as a site of entrenched interests. The same bodies that provide stability and legitimacy can also be exclusive in practice. The FASB's due process, for example, is shaped largely by the priorities of US-based public companies; the ICC's national committees are dominated by the trade ministries and industry associations of major economies; the IETF's rough consensus, for all its openness, still requires the time and technical fluency that participants from developing nations may struggle to contribute.

Resilience in a truly global cognitive infrastructure therefore requires more than preserving the human layer—it requires confronting the representational gaps within it. Without that breadth, the human layer does not merely risk creating standards that work for the few and fail to scale to the many; it risks forging a brittle form of resilience, one that is robust within a narrow set of dominant contexts but fragile when confronted with the full diversity of the global economies and communities it purports to serve. True resilience, then, lies not just in the durability of institutional processes but in their capacity to be legitimately claimed and shaped by all who depend on them.

The IETF and the IP Address Crisis (The 1990s)

By the early 1990s, the internet was growing at a pace that its own architects had not anticipated. The protocol that underpinned it—IPv4—contained a finite number of addresses, and exhaustion was no longer a distant theoretical concern but an imminent, mathematical certainty. The disruption on the horizon was nothing less than total network collapse: the system would simply run out of room.

The IETF's response exemplified the principle of designing for change, not for permanence. Rather than panic or attempt to mandate a single, comprehensive replacement overnight, the organisation leaned into its RFC process—a mechanism built precisely to accommodate evolution. This allowed for the deployment of stopgap measures such as NAT (Network Address Translation), which extended the life of IPv4, while simultaneously developing IPv6 as a long-term successor. Crucially, because the RFC framework treats standards as mutable documents that can be obsoleted incrementally rather than replaced through rupture, the

transition never forced a hard cutoff. The internet did not go offline; instead, it was quietly retrofitted while still carrying traffic. What appears in retrospect as a seamless evolution was in fact a deliberate institutional design that prioritised adaptability over the false promise of permanence.

The ICC and the Containerization Revolution (The 1960s–70s)

Before the standardized shipping container, international trade moved in sacks, crates, and barrels—each shipment a bespoke negotiation of risk and responsibility. The container changed everything, not merely in how goods were moved but in the legal architecture that governed trade. When a seller's obligation ended and a buyer's began—the core function of Incoterms—could no longer be defined by a ship's rail, the old rules became obsolete overnight.

The disruption was a technological revolution in logistics. The principle that proved decisive was preserving the human layer. The ICC did not allow the automated efficiency of cranes and container ships to dictate the law. Instead, its commissions—composed of the very traders, insurers, and lawyers who bore the risks—undertook the painstaking work of redefining terms. They introduced new frameworks like FCA (Free Carrier) to reflect a reality where goods were packed into sealed boxes miles from the port, with responsibility transferring at the moment of handover to a carrier, not at the ship's side. This human governance, deliberative and grounded in practice, prevented what would have been a cascade of insurance disputes and litigation. Yet it is worth noting that this same human layer, for all its expertise, has historically been shaped by the trade ministries and industry associations of major economies—a representational concentration that can make the resulting terms feel less like neutral standards and more like the defaults of established trading powers.

The FASB and the 2008 Financial Crisis

When the global financial system seized in 2008, accounting rules found themselves at the centre of the storm. "Fair value" accounting—the practice of marking assets to current market prices—was blamed for accelerating the crisis: as markets froze, asset values plummeted, forcing further write-downs that deepened the panic. The disruption was simultaneously economic and political, with legislators demanding immediate rule changes to obscure the losses.

The principle that held was building conflict resolution before you needed it. The FASB did not succumb to the pressure to abandon its standards arbitrarily. Instead, it activated the due process mechanisms it had spent decades constructing. Using its Basis for Conclusions as a foundation, it held emergency public hearings and issued clarifying guidance on how to value assets when markets were no longer active. The board did not hide the losses, but it provided a legitimate, transparent pathway for interpretation under extreme conditions. The result was that the financial reporting system retained credibility. Had the rules been changed overnight under political duress, the market's faith in the integrity of financial information would have evaporated. That the system survived is a testament to its procedural resilience—but a critical observer might also note that the FASB's due process, for all its legitimacy, has long been shaped

largely by the priorities of US-based public companies, leaving open the question of how well such mechanisms serve economies and stakeholders far from its traditional centre of gravity.

CONCLUSION

The Grammar of Global Commerce

The Anatomy of a Cognitive Infrastructure ends where it must: not with a blueprint but with a disposition. The ICC, IETF, and FASB/IASB are not bureaucratic obstacles to the efficiency promised by digital trade. They are the accumulated institutional wisdom of communities that have spent decades learning, at considerable cost, how to build rules that work. Their slowness is not inefficiency — it is the temporal signature of legitimacy.

The digitalisation of trade finance is genuinely transformative. The $2.5 trillion trade finance gap will not be closed by another round of banker-relationship management. It will be closed by digital identity systems that make the creditworthiness of a Kenyan smallholder legible to a Singaporean trade finance platform. It will be closed by tokenised invoices that convert illiquid receivables into marketable instruments accessible to global investors. It will be closed by agentic AI that eliminates the manual processing bottlenecks that make small transactions uneconomical for large banks.

But none of these transformations can occur without the legal infrastructure that the ICC is painstakingly building, the data infrastructure that the IETF is continuously improving, and the accounting infrastructure that the FASB and IASB are laboriously legitimising. The machine needs a legal framework before it can negotiate. The agent needs a protocol before it can communicate. The digital asset needs an accounting standard before it can appear on a balance sheet. Yet rules alone are not law. For digital trade to achieve its promise, the private rules written by institutions like the ICC must be anchored in a legal framework that grants them binding force across jurisdictions. That anchor is the UNCITRAL Model Law on Electronic Transferable Records (MLETR)—a UN-developed framework that countries including the United Kingdom, Singapore, and France are now adopting. Where the ICC provides the grammar of digital trade, the MLETR supplies the jurisdiction; together they transform electronic records from technical artefacts into instruments of legal obligation.

> *"Behind every flawless digital transaction, every trusted financial statement, and every container ship that reaches its destination lies a centuries-old human triumph: the ability to agree, not on everything, but on the essential rules that allow us to build, trade, and communicate together."*

The question for practitioners, regulators, and technology designers is not whether to embrace digital trade infrastructure. It is how to build systems that combine the efficiency of code with the legitimacy of law—that execute automatically but remain governable; that scale globally but remain contestable; and that move at the speed of software but remain accountable to the social processes that give rules their authority.

The invisible architecture is being rewritten. The quality of that rewriting will determine whether the next century of global trade is more equitable and more resilient than the last, or merely faster.

— End of Paper —

The greatest danger
in times of turbulence
is not the turbulence itself,
but to act with
yesterday's logic

— Peter Drucker

Discover more on:
www.owgconnect.com

Open Working Group