
ECONOMIC INTERESTS

INSTITUTIONS

CONTRACTS

SYSTEMS

AUTHORITY

INFORMATION

The Architecture of Institutional Credit

How the market evolved, why thirty years of modernization stalled, and what AI changes

A MONOGRAPH

The Architecture of Institutional Credit

How the market evolved, why thirty years of modernization stalled, and what AI changes

Steve Shyn

This paper is an attempt to explain the institutional credit market to a reader who has to make a decision about it.

It is not an investment recommendation, a legal opinion, or a compliance document.

CONEKTIQ

RESEARCH · SEPTEMBER 2026 · REVISED OCTOBER 2026

C O N E K T I Q

Published by Conektiq Research

2026

© 2026 Conektiq

conektiq.com

This paper is not an investment recommendation, a legal opinion, or a compliance document.

CONEKTIQ RESEARCH · SEPTEMBER 2026 · REVISED OCTOBER 2026

CONTENTS

Part One — Formation	<i>How the market developed</i>
Part Two — Anatomy	<i>How the participants fit together</i>
Part Three — Mechanism	<i>How the lifecycle actually works</i>
Part Four — Incumbency	<i>Why Loan IQ and ACBS won and stayed won</i>
Part Five — Standards	<i>What they solved and failed to solve</i>
Part Six — Technology around the incumbents	
Part Seven — Why collaboration repeatedly failed	
Part Eight — What private credit changes	
Part Nine — What AI changes, and what it does not	
Part Ten — Authority and truth	<i>In a multiparty credit market</i>
Part Eleven — Testing the hypothesis	
Part Twelve — What the industry could actually change	
Part Thirteen — Paths for a new entrant	
Appendices A–G	<i>including Appendix F — Collateral integrity</i>

A reader with limited time: Chapter 26, Chapter 41, Chapter 48, Chapter 68, Chapter 73, Chapter 81, Chapter 89. Those seven carry the argument.

DETAILED CONTENTS

Part One — Formation: how the market developed

1 The problem a syndicate solves · 2 From Eurodollar syndications to an asset class · 3 Institutionalization, 1995–1998 · 4 Collateralized loan obligations and the arrival of permanent demand · 5 The post-crisis re-plumbing · 6 The private credit era, and what it did to the operating problem · 7 What the history determines

Part Two — Anatomy: how the participants fit together

8 The sell side: two businesses inside one bank · 9 The agent bank as the market's registrar · 10 The buy side · 11 The service layer · 12 The infrastructure layer · 13 The consultancy and implementation layer · 14 Who pays whom · 15 A map of authority

Part Three — Mechanism: how the lifecycle actually works

16 Origination and primary syndication · 17 Closing, funding, and the birth of the register · 18 Servicing I: the notice as the unit of work · 19 Servicing II: interest, resets, fees, and the awkward arithmetic of SOFR · 20 Amendments, waivers, consents · 21 Secondary trading: price discovery without a tape · 22 Settlement: T+7, T+20, and what actually happens · 23 Delayed compensation and the problem of paying a moving syndicate · 24 Distress, restructuring, and the par/distressed boundary · 25 The duplication census · 26 Where the complexity actually lives

Part Four — Incumbency: why Loan IQ and ACBS won and stayed won

27 The core-system race · 28 Loan IQ: what is known, what is claimed, what is unknown · 29 ACBS and the other side of the field · 30 The anatomy of the moat · 31 The switching-cost arithmetic · 32 The dissatisfaction-switching gap · 33 The private-equity overlay · 34 The talent diaspora · 35 What the incumbency case does not prove

Part Five — Standards: what they solved and failed to solve

36 The documentation project · 37 The identifier project · 38 The messaging project: FpML and the anatomy of a failed standard · 39 The settlement project: twenty-five years · 40 The one that worked · 41 A theory of standardization failure in this market

Part Six — How technology evolved around the incumbents

42 Wave one: workflow digitization · 43 Wave two: data and pricing consolidation · 44 Wave three: settlement and reconciliation utilities · 45 Wave four: distributed ledger, and the graveyard · 46 Wave five: the bank consortia · 47 Wave six: artificial intelligence and document extraction · 48 The pattern across waves

Part Seven — Why collaboration repeatedly failed

49 The consortium as institutional form · 50 Coopetition: the documented rule · 51 The gatekeeper problem · 52 Symbiont: winning the lawsuit and dying anyway · 53 Who pays for transparency, and who loses from it

Part Eight — What private credit changes

54 Size, growth, and the arithmetic of the numbers · 55 Structurally different, not just smaller syndicates · 56 Valuation, opacity, and the arrival of the supervisor · 57 The operating layer: spreadsheets at scale · 58 The capital dynamic · 59 The stress test: First Brands, Tricolor, gating, and PIK · 60 Does private credit repeat the syndicated market's infrastructure history?

Part Nine — What AI changes, and what it does not

61 The 2022 baseline · 62 What is genuinely new · 63 The evaluation problem · 64 Why "AI reads contracts" is not a business · 65 What must remain deterministic · 66 The workforce collision · 67 The honest capability ledger

Part Ten — Authority and truth in a multiparty credit market

68 Who owns the truth about a loan? · 69 Authority by data element · 70 One golden source, or many with governed synchronization? · 71 Provenance versus consensus · 72 What verification can mean when the agent is contractually the truth · 73 Collateral: when legal authority does not establish economic fact · 74 The legal gate

Part Eleven — Testing the hypothesis

75 The coexistence hypothesis, stated falsifiably · 76 The case for · 77 The case against, in its strongest form · 78 The clock · 79 What would have to be true · 80 Paths the hypothesis does not consider · 81 Verdict on the central question

Part Twelve — What the industry could actually change

82 The distinction this part rests on · 83 What should remain bespoke · 84 What could usefully be standardized, and at which layer · 85 What should be translated rather than unified · 86 The agent-incentive problem, and the five ways to solve it · 87 Where regulation is actually required · 88 What AI should and should not be allowed to do · 89 The authority model the industry should accept · 90 What collaboration can and cannot solve · 91 What one institution can do alone, and a realistic sequence

Part Thirteen — Paths for a new entrant

92 Path A: the governed coexistence layer · 93 Path B: the private-credit administration platform · 94 Path C: the verification utility · 95 Path D: the channel play · 96 Path E: sell to the incumbents · 97 Path F: decline · 98 Decision criteria independent of path · 99 Closing

Appendices

A Chronology, 1970–2026 · B Participant register · C Claims register · D Figures this paper declines to use · E Open questions · F Collateral integrity · G Note on sources and method

A reader with limited time: Chapter 26, Chapter 41, Chapter 48, Chapter 68, Chapter 73, Chapter 81, Chapter 89. Those seven carry the argument.

LIST OF FIGURES

FIG. 1	Two markets, two architectures	<i>Chapter 1 · §1.1</i>
FIG. 2	One notice, many reconstructions	<i>Chapter 18 · §18.3</i>
FIG. 3	One economic event, many institutional states	<i>Chapter 25</i>
FIG. 4	The installed core is institution-specific	<i>Chapter 30 · §30.1</i>
FIG. 5	The one coordination that worked, and what it required	<i>Chapter 40</i>
FIG. 6	Where the cost falls, and where the benefit does	<i>Chapter 41</i>
FIG. 7	Six waves of modernization, and what survived them	<i>Chapter 48</i>
FIG. 8	What changed, and what did not	<i>Chapter 62</i>
FIG. 9	Analysis is not authority	<i>Chapter 64</i>
FIG. 10	The crossing between interpretation and execution	<i>Chapter 65 · §65.3</i>
FIG. 11	No participant sees the whole loan	<i>Chapter 68</i>
FIG. 12	Authority is allocated fact by fact	<i>Chapter 70</i>
FIG. 13	Legal certainty and economic fact are different questions	<i>Chapter 73</i>
FIG. 14	Many authoritative sources, governed synchronization	<i>Part Ten, closing</i>
FIG. 15	From proposition to institutional action	<i>Chapter 88</i>
FIG. 16	What one institution controls, and what it does not	<i>Chapter 91</i>

Sixteen numbered figures carry the argument's structural claims. Each entry links to its figure; the locator gives the chapter, since this document paginates on export. The paper's other exhibits — tables, maps and diagrams set within the chapters — are unnumbered by design.

PURPOSE

To examine how institutional credit evolved, why important operating problems persisted through decades of modernization, and what AI changes.

AUDIENCE

Anyone who works in, invests in, advises, builds for, regulates, or seeks to understand institutional credit.

STATUS

Independent research published by Conektiq Research.

IS NOT

A Conektiq sales document, product presentation, or argument for a predetermined technology solution.

The problem a syndicate solves

Begin with the thing itself.

A syndicated loan exists because a borrower wants more money than any single lender wants to lend it. That is the whole of the original problem. A corporation needs two billion dollars; no bank wishes to have two billion dollars of exposure to that corporation; so a group of banks lends the money together, on one set of terms, documented once.

Everything difficult about the market follows from the words *on one set of terms, documented once*.

Consider the alternative that was rejected. The borrower could have signed twenty separate bilateral loan agreements with twenty banks. Each would then be a normal, self-contained credit relationship: one lender, one borrower, one set of records, one set of decisions. The market did not do this, for good reasons. Twenty separate agreements means twenty separate negotiations, twenty sets of covenants that may conflict, twenty lenders who can each independently declare a default, and a borrower who must satisfy twenty different reporting regimes. The syndicated structure collapses all of that into one contract with one covenant package, one default definition, and one voting mechanism.

But collapsing twenty relationships into one contract does not collapse twenty lenders into one lender. The money still comes from twenty balance sheets and must be returned to twenty balance sheets. Somebody has to sit in the middle: receive the borrower's single payment and split it twenty ways, receive the borrower's single compliance certificate and distribute it twenty times, collect twenty votes when the borrower asks for a waiver, and — crucially — maintain the authoritative record of who the twenty lenders currently are.

That somebody is the **administrative agent**, and the agent is the load-bearing institution of this entire market. [FACT]

A note on terminology, because three words are used for one role and the paper will use all three. *Administrative agent* is the New York-law and LSTA term. *Facility agent* is the English-law and LMA term for substantially the same function. *Agent bank* is the market's ordinary speech, and it is imprecise in one respect worth knowing: the agent need not be a bank, and in private credit increasingly is not. Where the distinction matters this paper says so; elsewhere the three are interchangeable.

The agent is described in practitioner literature as "the central communication point": borrowers deliver financial statements and compliance materials to the agent, which distributes them to the lender syndicate; the agent maintains the **loan register**, the authoritative ledger of each lender's name and commitment; and the agent runs day-to-day servicing, payments, waivers and amendments. [FACT — DOCUMENTED

CONSISTENTLY ACROSS LOAN-AGENCY PRACTITIONER GUIDANCE] A facility agent's enumerated duties include arranging utilizations and interest-period selections, calculating and notifying interest, fees and break costs, operating accounts, receiving and distributing funds, and coordinating waivers, consents and amendments as the lenders' representative. [FACT]

Two consequences follow immediately, and they will recur in every part of this paper.

First: the agent is the market's designated authority on the register. Not by design and not by abuse, but by construction. The register is the contractual record of who owns the loan. There is only one register. The agent maintains it. Every other participant — every lender, every fund administrator, every custodian, every accountant — holds a copy or a derivative of the agent's record. When those records disagree, the agent's register governs operationally unless and until the discrepancy is successfully challenged or corrected. [FACT — THE REGISTER'S AUTHORITATIVE

STATUS IS ESTABLISHED IN CREDIT-AGREEMENT DOCUMENTATION AND CONFIRMED IN PRACTITIONER GUIDANCE] That is authority, not infallibility, and Part Ten returns to the distinction at length.

Second: the agent's work is a cost centre attached to a revenue centre. The bank that arranges the loan earns the arrangement fee, which is large. The same bank then usually becomes the agent, and the agency fee is small.

Post-close administration has been characterised as representing seventy to eighty per cent of the agent bank's operational burden over a facility's lifecycle. *[VENDOR CLAIM — THIS SPECIFIC PERCENTAGE APPEARS IN VENDOR*

COMMENTARY AND COULD NOT BE CORROBORATED INDEPENDENTLY; THE DIRECTIONAL POINT THAT SERVICING DOMINATES THE AGENT'S OPERATIONAL EFFORT IS WELL SUPPORTED] The economic asymmetry between the fee earned at the front and the work performed for years afterwards is the origin of nearly every operational pathology in Part Three.

1.1 Why the loan did not become a security

A reader from the equity or bond markets will already be asking the obvious question. Bonds solved this. A bond issue also raises money from many investors on one set of terms. Bonds are fungible, they carry standard identifiers, they settle in a day or two through a central depository, they trade on a reported tape, and nobody maintains a register of bondholders by hand.

Loans did not take that path, and the reason is legal rather than technological.

A syndicated loan is a **contract**, not a security. *[FACT]* It is not registered with the securities regulator, it is not reported to the corporate-bond transaction-reporting facility, and it is not held through a central securities depository. When a lender sells its position, it is not delivering a fungible instrument; it is assigning contractual rights and obligations, which requires — depending on the agreement — the consent of the borrower, the consent of the agent, or both. *[FACT]* Settlement therefore requires borrower

or administrative-agent consent to assignment, negotiated transfer documentation, checks against borrower-maintained "disqualified lender" lists, and know-your-customer diligence on each new lender; and there is no central securities depository and no netting facility of the kind that clears equities. *[FACT]*

This is worth dwelling on because it is frequently misdescribed. The loan market is not slow because its participants are backward. It is slow because a loan trade is an **assignment of contractual rights and obligations between identified parties**, and the effectiveness of that transfer depends on the consent and other conditions the parties negotiated into the agreement.

The bond market and the loan market therefore modernized around different architectures. Securities markets progressively concentrated functions into common rails: standardized instruments, identifiers, depositories, transaction reporting and increasingly electronic execution. The loan market largely built interfaces between institutions that continued to perform different contractual roles and maintain different records.

Put differently: the bond market could concentrate more of its operational complexity. The loan market distributed much of its complexity among the parties to the contract. *[INFERENCE]*

That does not mean bond markets are simple, or that every feature of a loan could not be standardized. It means that making a loan behave operationally like a bond eventually encounters the legal form of the instrument itself. Chapter 26 formalises the point. For now, the important conclusion is narrower: some of the market's most persistent operating friction was determined at the outset by choices of legal form and institutional role, and software alone cannot repeal those choices.


FIGURE 1

Two markets, two architectures

PANEL A

Corporate bond market

Standardized instrument. Centralized infrastructure. Many participants, convergent architecture.



Around the rails: regulators (SEC, FINRA) for trade reporting and surveillance; data providers for pricing and evaluated prices; service providers for corporate actions, income and tax.

Key characteristics. Standardized instrument terms · centralized clearing and settlement · high straight-through processing · agreed economics settled through shared infrastructure · complexity exists, but is concentrated in common infrastructure.

Result: many participants, fewer independent states. The market built rails.


The bond market concentrated its operational complexity into common rails — standardized instruments, identifiers, a depository, transaction reporting. The loan market built interfaces between institutions that continued to perform different contractual roles and maintain different records. §1.1 and Chapter 26 set out why that difference follows from legal form rather than from technical capability. [INFERENCE]

FIGURE 1, CONTINUED

PANEL B

Syndicated loan market

Contract-specific terms. Distributed infrastructure. Many participants, many independent states.



Ongoing lifecycle. Payments · accruals · covenant compliance · amendments · consents · waivers · notices · assignments · reconciliations · investor reporting · regulator inquiries · secondary trades · portfolio transfers · asset sales · wind-down.

Key characteristics. Contract-specific economics and documentation · authority distributed by fact — different facts are authoritative in different places, and the agent register is authoritative for lender-of-record state · no universal source authoritative for every economically relevant fact · multiple legitimate representations of the same event · heavy reliance on manual processes, notices and reconciliations · no centralized settlement · complexity is structural, not a software problem.

Result: many participants, many independent states. The market built interfaces.

SOME QUESTIONS HAVE NO SINGLE AUTHORITATIVE SOURCE

What is the "true" market price? · Has this loan been pledged twice? · Are all consents valid? · Is every lender record correct? · Is all relevant evidence complete?

Sources: DTCC, ISDA, LSTA, FSB and industry reporting. Illustrative and simplified.

From Eurodollar syndications to an asset class

The syndicated loan as a technique is old. The syndicated loan as an *asset class* — something an investor buys rather than something a bank originates — is not.

Through the 1970s, syndicated lending was primarily a mechanism for banks to share exposure among themselves: sovereign lending, project finance, large corporate facilities funded out of the Eurodollar market. The participants were banks, the loans stayed on bank balance sheets, and the concept of a secondary market barely existed. A bank that wanted out of a loan asked another bank to take it, as a favour, and the paperwork was whatever the two banks agreed.

Two developments in the 1980s changed this.

The first was the leveraged buyout. Highly leveraged corporate acquisitions required more debt than the acquirer's cash flows comfortably supported, at pricing that reflected the risk. That produced a class of loan that behaved less like relationship banking and more like a high-yield security: priced off credit risk rather than relationship, held for return rather than for the customer relationship, and — critically — attractive to investors who were not banks.

The second was distress. Loans made in the early 1980s, both to sovereign borrowers and to leveraged corporates, deteriorated. Banks wanted to sell them. Buyers emerged who specialised in buying distressed claims. A secondary market began to function, in a rudimentary way, out of necessity rather than design.

By the early 1990s a recognisable market structure existed: banks originated and syndicated loans; some of those loans traded; a set of specialist trading desks made markets in them; and a growing

population of non-bank investors bought them. What did not exist was any standardisation of how any of it was documented, priced, settled, or recorded.

This is the pivotal moment for understanding everything that follows.

[INFERENCE] The market industrialised — in volume, in participant count, in trading activity — *before* it standardised. By the time anyone attempted to impose common documentation, common identifiers, or common settlement conventions, there were already thousands of live agreements, dozens of institutions with established internal processes, and a set of intermediaries whose economics depended on the existing arrangements. Standardisation in this market has never been a green-field design exercise. It has always been a retrofit onto a functioning system with entrenched interests, which is why Part Five reads the way it does.

Institutionalization, 1995–1998

Between 1995 and 1998 the market acquired its institutions.

The **Loan Syndications and Trading Association** was founded in 1995 in New York. The **Loan Market Association** was founded in London in December 1996. The **Asia Pacific Loan Market Association** was founded in 1998 in Hong Kong by fifteen major international banks. *[FACT]* All three were created for substantially the same purpose: to develop best practices and standard documentation in order to grow liquidity in secondary trading, and all three describe themselves as working toward that shared aim. *[FACT]* The LMA today reports more than 880 members across 69 jurisdictions; the APLMA reports more than 400. *[FACT]*

The founding purpose is worth reading carefully, because it explains both what these bodies achieved and what they could not.

They were founded **to grow liquidity**, and they were founded **by market participants**. Neither body is a regulator. Neither can compel anything. Their instruments are model documents, recommended conventions, published timelines, and the soft power of being the place where the market's practice is codified. That is a genuine and underrated form of authority — but it is authority over *convention*, not over *behaviour*, and the difference matters enormously when the convention costs somebody money.

The three bodies did converge on some things. They jointly issued the Green Loan Principles in 2018 and the Sustainability-Linked Loan Principles in 2019, with jointly updated guidance in February 2023, and jointly launched a standard Administrative Details Form for global use. *[FACT]* These are real accomplishments and they show that cross-body coordination is possible when the subject is additive — a new labelling standard for a new product category, imposed on nobody who does not want it.

They did not converge on the things that would have mattered most. LSTA documents govern New York-law credit agreements; LMA documents govern English and European-law agreements; APLMA documents govern Asian practice. The bodies produced a comparative users' guide and a shared glossary "with the aim of eliminating confusion" — which is, read plainly, an acknowledgment that two distinct frameworks persist and will continue to. [FACT] Even a deliberate convergence attempt — a new LSTA form of revolving credit facility — has been analysed by practitioners as still notably different from the LMA equivalent. [FACT] Cross-border document mismatch is an active legal risk, examined in law-firm client alerts on the hazards of trading English-law loans on New York-law par terms. [FACT] And APLMA documentation is characterised as using term-sheet form with greater optionality and flexibility, tuned to Asian market practice, rather than adopting LMA or LSTA templates — meaning the Asian market has *less* documentary convergence, not more. [FACT]

Chapter 36 examines why. The short version, which will be developed at length, is that model documentation standardises boilerplate and cannot standardise economics. [FACT] The LSTA's model provisions reduce negotiation on procedural terms, but "the provisions that have more commercial effect on the parties (such as mandatory prepayments, business undertakings, financial covenants, representations and warranties, conditions to drawdown, etc.) remain as bespoke to the specific transaction as ever," and the recommended standard forms "are rarely used as a starting draft for negotiation." [FACT]

That sentence — *rarely used as a starting draft* — should be read twice by anyone who imagines that thirty years of standard-setting produced standard documents.

Collateralized loan obligations and the arrival of permanent demand

The single most important change in the market's demand structure was the collateralized loan obligation.

A CLO is a vehicle that buys a portfolio of leveraged loans and funds itself by issuing tranchised securities against them. Its significance for this paper is not its credit structure but its behaviour: a CLO is a permanent, mandated, non-discretionary buyer of leveraged loans, with a defined reinvestment period and compliance tests that require it to hold a portfolio meeting specified criteria. It does not decide each morning whether it feels like being in the loan market.

The scale of what this created is difficult to overstate. CLOs purchased approximately sixty-one per cent of all new-issue leveraged loans in 2024 and own roughly sixty-four per cent of the leveraged loan market outright, with mutual funds at roughly nineteen per cent, banks at roughly eight and insurance companies at roughly six. *[FACT — AS ORDERS OF MAGNITUDE. THESE SPLITS COME FROM A CONGRESSIONAL RESEARCH PRODUCT REACHED THROUGH A SEARCH SUMMARY AFTER THE PRIMARY DOCUMENT REFUSED AUTOMATED RETRIEVAL, CORROBORATED BY AN ASSET MANAGER'S INVESTOR PRIMER. THEY SHOULD BE READ AS APPROXIMATE SHARES, NOT TO THE DECIMAL.]* Independent estimates place CLO holdings at sixty-five to seventy per cent of the broadly syndicated loan market and CLO demand at fifty to seventy-five per cent of newly issued highly leveraged loans. *[VENDOR CLAIM]*

Sizing the CLO market requires care, because the available estimates are not comparable and cannot be added. The commonly cited figures are: global CLOs outstanding of approximately \$1.4 trillion as of April 2025, with one source putting the same measure nearer \$1.2 trillion; US CLOs outstanding of approximately \$1.5 trillion; and European CLOs of approximately €300 billion, the largest single segment of European

securitisation at more than forty per cent. *[FACT — EACH FIGURE IS SEPARATELY REPORTED]* The reader will notice that the US figure exceeds the global one. That is not a transcription error in this paper; it is a definitional collision in the sources, which differ on whether "CLO market" means notes outstanding, collateral under management, or aggregate deal size, and on as-of dates spread across a year in which issuance and repayment were both running at records. **The honest statement is that the global market is of the order of \$1.2–1.5 trillion and the US is the large majority of it; anyone who needs a precise number needs to specify the measure first.** *[INFERENCE]*

A separate definitional warning applies throughout this paper. **"Broadly syndicated loan," abbreviated BSL, means a large loan underwritten and administered by commercial or investment banks, rated by agencies, sold to a wide investor base, and actively traded in the secondary market.** *[FACT]* It is the segment CLOs buy. It is distinct from club deals, from bilateral loans, and from private credit, and the paper keeps them separate because their operating models differ (Chapter 55).

The consequence for the operating problem is direct and is often overlooked. **The dominant owner of the leveraged loan market is a vehicle that must, by contract, produce daily portfolio compliance data, monthly trustee reports, quarterly waterfall calculations, and audited investor reporting — on assets whose authoritative records are held by dozens of different agent banks and delivered by unstructured notice.** *[INFERENCE]* The CLO did not merely add demand. It added a demand for *data* of a kind the market's plumbing was never built to supply, and it added that demand at industrial scale.

The market's regulatory arrangements amplified this. In February 2018, the DC Circuit held in *LSTA v. SEC* that managers of open-market CLOs are not "securitizers" within the meaning of the Dodd-Frank risk-retention provisions and are therefore exempt from the five per cent retention requirement — reasoning that an open-market manager never holds,

transfers or retains the loans and so cannot "retain" risk. *[FACT]* The decision reversed a 2016 district-court ruling. Market commentary widely credits it with unlocking the subsequent growth in US CLO issuance — an attribution this paper reports rather than endorses, since the same period saw rate, spread and demand conditions that would independently have supported issuance. *[FACT THAT THE ATTRIBUTION IS WIDELY MADE; INFERENCE AS TO HOW MUCH OF THE GROWTH IT EXPLAINS]* European and UK CLOs remain subject to a five per cent retention rule, which is one of several reasons the two markets have diverged. *[FACT]*

By 2024–2025 the CLO machine was running at record output: US CLO new issuance of roughly \$202 billion in 2024 and roughly \$209 billion in 2025, alongside approximately \$232 billion of resets and \$105 billion of refinancings in 2025. *[FACT]* Private-credit CLOs reached roughly \$42 billion in 2024, about a fifth of CLO issuance — a number worth remembering when Part Eight considers whether private credit is diverging from or converging on the syndicated market's structure. *[FACT]*

The post-crisis re-plumbing

The financial crisis and its regulatory aftermath did something to this market that is usually described in terms of capital ratios but is better understood in terms of *where credit is manufactured*.

Three regulatory changes mattered. Basel III raised the capital cost of holding leveraged corporate exposures on a bank balance sheet. The Volcker Rule constrained certain bank activities. And in 2013 the Federal Reserve, the Office of the Comptroller of the Currency and the Federal Deposit Insurance Corporation issued interagency guidance on leveraged lending, which subjected transactions above roughly six times total debt to EBITDA to heightened supervisory scrutiny and required banks to articulate risk appetite and conduct stress testing. *[FACT]*

The stated intention was to make bank leveraged lending safer. What happened next is usually described as though the guidance caused the growth of private credit, and the honest position is more careful than that — though in this instance the causal evidence is unusually good, because it comes from the regulators themselves.

When the OCC and FDIC rescinded the guidance in December 2025 they stated that it had pushed leveraged lending outside the regulatory perimeter and into private credit. *[FACT]* That is a supervisory admission of the mechanism, not an outside inference. Independently, the Bank for International Settlements has framed the shift as a change in the *structure* of credit intermediation, from bank to non-bank, rather than a change in the aggregate level of credit. *[FACT]*

Those two together support a reasonably strong claim: the regulatory package raised the relative cost of holding leveraged exposures on a bank balance sheet, and lending relocated rather than disappeared. *[INFERENCE,*

What they do not support is a claim that regulation was the *only* driver, and the next paragraph gives the reason.

The relocation was substantial. Private credit expanded from \$158 billion globally in 2010 to close to \$2 trillion by mid-2024, according to the International Monetary Fund. [FACT] The Federal Reserve, using a more conservative definition, put the market at roughly \$1.34 trillion in the United States and roughly \$2 trillion globally as of the second quarter of 2024, having grown roughly fivefold since 2009. [FACT] Business development companies — the public window into private credit — grew from roughly \$30 billion in 2010 to roughly \$450 billion of fair value and \$475–500 billion of assets under management by the first quarter of 2025, with the top ten BDCs accounting for fifty-three per cent of the market. [FACT]

The borrower-side pull was as real as the regulatory push: speed of execution, certainty of close, confidentiality, and flexibility of structure, relative to a syndication that requires a ratings process, a marketing period and a book-build. [FACT]

5.1 The reversal

Then, in December 2025, the regulators changed their minds.

On 5 December 2025 the OCC and the FDIC formally rescinded the 2013 interagency leveraged lending guidance and the associated 2014 frequently-asked-questions document, describing the guidance as "overly restrictive" and explicitly acknowledging that it had pushed leveraged lending outside the regulatory perimeter and into private credit. [FACT] The Federal Reserve, a co-author of the original guidance, had not joined the rescission as of the date of this paper. [FACT] The legal groundwork had been laid years earlier: a 2017 Government Accountability Office ruling held that the guidance was a "rule" under the Congressional Review Act that had never been submitted to Congress. [FACT]

This is one of the more consequential and least discussed facts in the current landscape, and it cuts in two directions at once.

It partially re-opens the bank channel. Banks that had ceded leveraged lending to private credit funds because holding it was expensive may now find it less so. *[INFERENCE]* More loan volume flows through bank balance sheets and therefore through the agent-bank servicing machinery.

But it does not reverse the migration, because the migration was not purely regulatory. Private credit's borrower-side advantages — speed, certainty, confidentiality, flexibility — are structural rather than regulatory, and the largest managers have by now built origination franchises and permanent capital bases that do not disappear because a supervisory document was withdrawn. *[INFERENCE]*

The most likely outcome is neither reversal nor continuation but **hybridisation**: more loan volume flowing through joint structures in which a bank originates or distributes and a manager holds, which is precisely what the 2024–2026 wave of bank–manager partnerships suggests. Part Eight documents that wave. *[INFERENCE]*

The private credit era, and what it did to the operating problem

Private credit's arrival is usually discussed as a credit story: are the loans good, is the leverage layered, are the valuations honest. Part Eight takes those questions up. This chapter is concerned with something narrower — what private credit did to the *operating* problem.

The answer is that it made it worse, in a specific and instructive way.

Broadly syndicated lending, whatever its inefficiencies, has a market infrastructure. It has agent banks with servicing platforms. It has trade-body documentation, however imperfectly converged. It has settlement platforms, ratings, pricing services, identifiers, and a CLO machine with defined reporting obligations. It is slow, manual and fragmented, but it is *institutionalised*.

Private credit is bilateral or club-originated, hold-to-maturity, floating-rate, and — the point that matters here — valued by the manager rather than marked to a market. *[FACT]* It has no agent bank in the syndicated sense for a large share of transactions, no trade-body documentation with the reach of LSTA or LMA standards, no meaningful secondary market for most positions, and no equivalent of the settlement rails.

What it has instead is spreadsheets. The characterisation across the industry's own literature is consistent: private credit is "built around bilateral lending relationships, tailored documentation, and buy-and-hold structures," where "each deal comes with unique covenants, payment schedules, and waterfall structures," and data arrives as "PDFs, emails and spreadsheets... each structured differently and requiring manual reading."

[VENDOR CLAIM] "A spreadsheet that worked for 20 loans breaks down under the weight of 200." *[VENDOR CLAIM]* "Normalising unstructured information — credit agreements, agent bank notices, portfolio company financials — remains the sector's biggest operational challenge." *[VENDOR CLAIM]*

Those are vendor characterisations and should be read as such. But the Financial Stability Board — which sells nothing — independently names "data challenges faced by regulators when monitoring exposures" among its four vulnerability clusters in private credit, and the Bank of England has stated that data gaps mean banks "cannot systematically aggregate... combined credit and counterparty risk exposures." *[FACT]* When the vendors and the supervisors describe the same gap, the gap is probably there.

So the market's centre of gravity has been moving, over fifteen years, from a segment that is badly served by its infrastructure toward a segment that has almost none. *[INFERENCE]* Whatever one concludes about the opportunity, that is the trend line.

What the history determines

Five structural facts were fixed by the history in this part, and none of them can be changed by a vendor.

One: the loan is a contract, not a security. This determines that transfer requires consent, that there is no central depository, that settlement is a documentary process rather than a book entry, and that no amount of electronification produces bond-like settlement without changing the instrument. *[FACT]*

Two: the agent holds the register, and the register governs. This determines that every other participant's ownership record is derivative, that reconciliation is structurally necessary rather than a failure of diligence, and that any system seeking authority over lender-of-record state must either be the agent, be authorised by the agent, or accept a subordinate role. *[FACT]* The register can contain an error; its significance is that it remains the operative record until corrected.

Three: the market industrialised before it standardised. This determines that standardisation is always a retrofit, always faces entrenched interests, and always requires somebody to bear a cost for a benefit that accrues elsewhere. *[INFERENCE]*

Four: economics standardised least. Model documents converged on boilerplate and diverged on covenants, pricing grids and prepayment mechanics — which is to say, on precisely the terms a servicing system must model. *[FACT]*

Five: the buy side became the dominant owner without becoming the dominant authority. CLOs own roughly sixty-four per cent of the leveraged loan market and control none of its data. *[FACT]* The party with the greatest need for timely, accurate position data is structurally the party furthest from its source.

Taken together, these facts imply one further proposition that Part Two begins to map and Part Ten will establish precisely: there is no single institution that is authoritative for every economically relevant fact about a loan. The agent's register governs lender-of-record state; the borrower is authoritative for specified information about its own condition; lenders, custodians, trustees, administrators and other parties maintain records for functions over which they have their own responsibilities and, in some cases, their own authority. *[INFERENCE]*

The market's fragmentation is therefore not entirely accidental. Some of it is the operational consequence of a contractual structure that deliberately allocates different responsibilities to different parties.

Everything in Parts Two through Seven is a consequence of these five facts. Everything in Parts Eleven and Twelve must be tested against them.

II

PART TWO

Anatomy

How the participants fit together

A market is not a list of firms. It is a set of positions, each defined by what the occupant controls, what it is paid for, and what it must obtain from someone else. This part maps those positions.

The mapping matters more here than in most markets, because the institutional credit market's most important structural feature is that **the party with the data is not the party that needs it**. Understanding who sits where is therefore not background; it is the analysis.

The sell side: two businesses inside one bank

ORIGINATION & DISTRIBUTION

Profit centre. Arrangement fee at the front. Measured on league-table position. Relationship to the transaction substantially ends at closing.

AGENCY & SERVICING

Cost centre. Small annual fee, set privately. Runs for the life of the facility — up to seven years.

The phrase "the banks" obscures the fact that a large arranger bank contains at least two businesses with different economics, different time horizons, and different attitudes to the operating problem.

The first is origination and distribution. A syndicate desk wins a mandate, structures a facility, underwrites or best-efforts it, markets it to investors, builds a book, allocates, and closes. It is paid an arrangement fee at the front, it is measured on league-table position and fee revenue, and its relationship with a given transaction substantially ends at closing.

The second is agency and servicing. Once the loan closes, someone must run it: maintain the register, calculate interest, issue notices, process drawdowns and rollovers, tabulate amendment votes, process assignments, and distribute cash — for the life of the facility, which may be seven years. This function earns an annual agency fee that is set privately in a fee letter and is, by every available account, small relative to the arrangement fee.

[FACT — THAT AGENCY FEES ARE SET IN PRIVATE FEE LETTERS AND NOT PUBLICLY RANGED IS DOCUMENTED; THE RELATIVE MAGNITUDE IS AN INFERENCE FROM THE ABSENCE OF ANY PUBLIC CLAIM TO THE CONTRARY AND FROM THE TREATMENT OF AGENCY AS A COST CENTRE IN THE OPERATIONAL LITERATURE] [INFERENCE]

This split is the origin of a great deal.

The origination business is a profit centre with visible revenue attribution and executive attention. The servicing business is a cost centre. Senior loan-operations heads at leading banks are described as indicating that "fundamental issues in loan operations receive only surface-level attention," with banks prioritising customer acquisition, and that risk, efficiency and process are "not on the leadership agenda at most banks."

[VENDOR CLAIM — FROM A CONSULTANCY THAT SELLS INTO THIS PROBLEM; DIRECTIONALLY CORROBORATED BY THE OFFSHORING AND RELOCATION BEHAVIOUR DOCUMENTED BELOW] [INFERENCE]

The consequence is that the function which holds the market's authoritative data is the function with the least claim on the institution's investment budget. That is not an accident of any particular bank's management. It is what happens when a fee is earned at the front and the work is performed at the back. *[INFERENCE]*

8.1 The relocation of the back office

The clearest evidence for the cost-centre characterisation is not what banks say but where they put the work.

In 2019 JPMorgan was reported to be weighing shifting thousands of New York-area jobs to lower-cost hubs — Plano, Texas; Columbus, Ohio; Wilmington, Delaware — explicitly targeting "non-client-facing... back and middle-office professionals... in compliance, operations, trade support, accounting and human resources." *[FACT]* By 2024 JPMorgan employed roughly 31,500 people in Texas, its largest state presence, surpassing New York, including approximately 25,000 in Plano. *[FACT]* Goldman Sachs ran a cost initiative that asked selected New York managers to relocate to Salt Lake City or Dallas "or leave the company"; the growth of those offices has been "fueled by junior-ranking employees, often in middle- and back-office functions," and Salt Lake City has hosted Goldman's US back office since 2000 and is now its third-largest North American site. *[FACT]* Bank of America and Citi operate large operations centres in Tampa and Jacksonville. *[FACT]*

Delaware's role has a specific origin. The Financial Center Development Act of 1981 eliminated interest-rate ceilings on revolving and closed-end credit, permitted card fees, and offered declining bank tax rates as income rose, with a requirement that out-of-state holding companies employ at least a hundred people within a year. *[FACT]* By 1991 Delaware hosted thirty-one credit-card and wholesale banks with roughly 13,000 direct and 7,000 related jobs. *[FACT]* JPMorgan Chase is Delaware's largest private employer with approximately 10,000–11,000 staff in northern Delaware. *[FACT]*

Beyond the domestic hubs is the offshore layer. India's global capability centre base is reported at roughly 2,100 centres and 2.4 million professionals; within that, financial services is a leading vertical, with reported figures of more than 300 captive centres in banking, financial services and insurance exceeding 600,000 staff, and more than ninety global financial institutions operating some 185 centres. *[VENDOR CLAIM — THESE COUNTS REACH THE PUBLIC RECORD THROUGH INDUSTRY AGGREGATORS RATHER THAN PRIMARY DISCLOSURE, AND THE RESEARCH PASS THAT ASSEMBLED THEM FLAGGED THEM FOR VERIFICATION AT SOURCE. THEY ESTABLISH AN ORDER OF MAGNITUDE AND SHOULD NOT BE QUOTED AS MEASUREMENTS.]* Business-process outsourcers run lending operations at scale: one major provider states it "handles more than \$200 billion in assets" across commercial-lending processes "from origination to servicing." *[VENDOR CLAIM]* Typical offshore cost savings in financial services are cited at forty to sixty per cent versus onshore. *[VENDOR CLAIM]*

There is a directly relevant caveat about what that cost saving buys. Industry research on financial-services outsourcing reports first-year attrition of up to roughly fifty per cent in traditional offshore centres, with the observation that "after eighteen months the team that understood your products and regulatory obligations has entirely turned over." *[VENDOR CLAIM — FROM FIRMS SELLING OUTSOURCING SERVICES, AND THEREFORE NOT DISINTERESTED; THE DIRECTION IS NONETHELESS UNFLATTERING TO THE SELLER]* Consultancy research reports several banks facing attrition rates of twenty to seventy per cent among onshore, US-centric loan-operations talent, "leading to a substantial loss of institutional knowledge," against a time to productivity for new hires

of one to two years. *[VENDOR CLAIM — THE CONSULTANCY SELLS MODERNIZATION SERVICES, THE PRIMARY DOCUMENT COULD NOT BE RETRIEVED DIRECTLY DURING RESEARCH, AND THE FIGURES ARE QUOTED FROM SEARCH EXTRACTS. THEY ARE THE MOST-CITED NUMBERS IN THIS FIELD AND AMONG THE LEAST VERIFIED.]*

Chapter 66 examines what this means for the market's human capability. Here the point is structural, and it should be stated as a description rather than an accusation: the function holding the authoritative record of a multi-trillion-dollar market has been systematically relocated toward lower-cost labour markets with higher turnover, running processes that remain substantially manual. That is the rational response of an institution to a cost centre, not a failure of management. Whether it *causes* operational failures is a separate question, and Chapter 8.2 examines the one case where the public record is detailed enough to say anything at all.

8.2 Revlon

On 11 August 2020, Citibank, acting as administrative agent on a 2016 syndicated term loan owed by Revlon, intended to remit an accrued interest payment of approximately \$7.8 million to lenders. It instead wired approximately \$893.9 million — the loan's entire outstanding principal — on a facility not due until 2023. *[FACT]*

The mechanism deserves to be stated exactly, because it is the single best-documented illustration of how this market's operating model actually fails.

The payment ran through Citi's "six-eye" three-person approval procedure. A "maker" — a contractor employee — inputs the instruction into the loan-processing system; a "checker," a second contractor employee, verifies it; an "approver," a Citibank senior manager, gives final sign-off. To route only the interest and suppress the principal, the operator needed to direct the principal to an internal "wash account" by setting three fields: "PRINCIPAL," "FRONT," and "FUND." The operator believed that ticking the "PRINCIPAL" checkbox plus entering the wash-account number was sufficient. Because

"FRONT" and "FUND" were not also set to the wash account, the full principal was released to the external lenders. [FACT] The control team consisted of six Citibank employees based in Delaware and nine employees of an outsourcing provider in India, with the final-say supervisor in Delaware. [FACT]

Some lenders voluntarily returned roughly \$385 million. Ten asset managers holding approximately \$500 million refused, and litigation followed. In February 2021 the Southern District of New York ruled that the non-returning lenders could keep the money under New York's discharge-for-value doctrine, the judge describing the incident as "one of the biggest blunders in banking history." [FACT] In September 2022 the Second Circuit vacated that ruling, holding that the discharge-for-value defence applies only where the debt is presently due, and that because the loan was not due until 2023 the defence was unavailable. [FACT] Rehearing en banc was denied in October 2022; the parties settled in December 2022. [FACT] There is no evidence of any Supreme Court involvement. [FACT]

Separately, in October 2020, the OCC assessed a \$400 million civil money penalty against Citibank for deficiencies in enterprise-wide risk management, data governance and internal controls, including risk-data aggregation and regulatory reporting, and issued a cease-and-desist order; the Federal Reserve took a related action against the holding company. In 2024 regulators fined Citi a further reported \$136 million for insufficient progress on the 2020 consent order. [FACT]

The pattern recurred. In April 2024 Citigroup credited a customer's account with \$81 trillion instead of \$280, after a payments employee failed to delete pre-filled zeros in a backup system's manual-entry field; the error was missed by both the maker and the checking official and caught by a third employee roughly ninety minutes later. It was reversed before funds left the bank and disclosed to the Federal Reserve and the OCC as a near miss. [FACT]

The pattern is not confined to one institution. In March 2018 Deutsche Bank inadvertently transferred €28 billion — roughly \$35 billion, more than its own market capitalisation — to its account at an exchange during a routine daily collateral adjustment, because euros were input instead of yen; the error was caught within minutes and reversed with no loss. *[FACT]* In June 2015 a junior foreign-exchange trader at the same bank paid approximately \$6 billion to a hedge-fund client by processing a gross figure instead of a net amount; the funds were recovered the next day and the incident reported to three regulators. *[FACT]*

Four incidents, three institutions, one failure mode: **a human being enters a number into a screen, other human beings look at the screen, and the screen does not make the error visible.** *[INFERENCE]* Adding reviewers to a manual process does not remove the tail risk; it multiplies the number of people who can fail to see the same thing. *[INFERENCE]*

What this case does and does not establish about offshoring. The composition of the control chain — six staff in a low-cost domestic hub, nine at an offshore provider, the approving supervisor in the domestic hub — is documented. *[FACT]* That the chain failed is documented. That the *first* fact caused the second is not: the same institution produced the same failure mode four years later in a different process, and two of the four incidents above occurred at an institution whose relevant desks were in London and Frankfurt. The defensible reading is that **the failure mode is a property of manual maker-checker control over a confusing interface, and that relocation and turnover plausibly raise its frequency without being its cause.** *[INFERENCE]* Readers who want the stronger claim — that offshoring caused this — will not find it supported here, and the practitioner commentary that gestures at it in the sources is opinion rather than finding.

The agent bank as the market's registrar

The agent's position deserves a chapter of its own because it is the position from which everything else in this market is negotiated.

The agent's functional duties have already been enumerated: distribute notices, oversee timely and accurate payments, coordinate assignments and lender onboarding, maintain records, calculate and notify interest, fees and break costs, operate accounts, receive and distribute funds, and coordinate waivers, consents and amendments. *[FACT]* The agent maintains the register — the ledger of lender names, commitments and principal amounts — updating it from each assignment and acceptance, and a buyer must fund upon the agent's notice that a trade is effective in the register. *[FACT]*

Three properties of this position are load-bearing for the rest of the paper.

The agent is the origin of the data. Notices to lenders are generated by the agent's core servicing system. Loan IQ, the dominant such system, has modules explicitly covering "bills / notices" and "enhanced notice generation." *[VENDOR CLAIM]* Any effort to modernise the market's data flow must therefore either capture data at that origin, or normalise the notices after they leave it — and those are the two strategies that have actually been attempted. Chapter 44 and Chapter 46 examine both. *[INFERENCE]*

The agent is usually the arranger. The bank that wins the mandate typically takes the agency, which means that the institutions holding the market's data are the same institutions that dominate its origination. *[FACT]* The top-of-market concentration figures in Chapter 14 are therefore also concentration figures for control of loan data.

The agent has discretion, and the discretion is priced. When a lender sells its position, the agent charges a processing fee — \$3,500 under the standard form, with some agreements setting \$4,500 — payable in immediately

available funds and waivable "in the Administrative Agent's sole discretion." [FACT] The agent both sets and collects this fee and controls whether the trade and its consents move.

The structural conflict is documented rather than alleged. Practitioner literature notes that an entity acting as both lender and agent "can find itself conflicted between acting in the best interests of its own institution versus the wider lender group," and describes running the agency desk as "an additional income fee stream." [FACT] Buy-side complaints about agent processing are long-standing and specific: administrative agents "often do not have the capacity to review every trade within one business day," with review and acknowledgement delays "lasting three or more days," compounded by tax validations, borrower-consent requirements, and periodic trading "holds" during which agents will not process assignments. [FACT] And the delayed-compensation mechanism — the market's remedy for late settlement — explicitly does not cover delays caused by failure to obtain agreed consents, leaving the agent and borrower consent chokepoint without a remedy. [FACT]

It is important to state what this evidence does and does not establish. It establishes a structural conflict, a discretionary fee, and persistent documented delay complaints. It does not establish deliberate extraction or weaponisation of consent; no regulator has made such a finding, and this paper does not assert one. [INFERENCE] The honest characterisation is that the agent occupies a position in which its interests and the buy side's interests diverge, that the market has noticed, and that the market's own remedy conspicuously exempts the part of the process the agent controls.

The buy side

The buy side is where the reconciliation burden lands, and it is worth being precise about who is carrying it.

CLO managers are the largest holder class, buying sixty-one per cent of new-issue leveraged loans in 2024 and owning roughly sixty-four per cent of the market. *[FACT]* Their operational obligations are the heaviest in the market: daily portfolio compliance testing, waterfall calculations, trustee reporting, and investor reporting, all against positions whose authoritative records sit at agent banks.

Loan mutual funds and exchange-traded funds, including a fast-growing CLO ETF segment that added roughly \$16 billion to reach \$22.5 billion of assets in 2024. *[FACT]* These carry daily net-asset-value obligations against an asset class that settles in weeks.

Business development companies and private-credit funds. BDC assets reached roughly \$450 billion of fair value by the first quarter of 2025, with the top ten representing fifty-three per cent, and perpetual non-traded BDCs the fastest-growing vehicle. *[FACT]* A 2025 asset-manager survey reported that ninety-four per cent of institutional investors now hold private credit. *[FACT THAT THE SURVEY REPORTED IT; THE FIGURE IS A SURVEY RESULT RATHER THAN A CENSUS, AND THE SURVEYING FIRM SELLS INTO THE ASSET CLASS]*

Insurance companies, holding roughly six per cent of leveraged loans directly and buying CLO tranches, with statutory, GAAP and tax reporting obligations that are materially more demanding than a fund's. *[FACT]* Private credit on US life-insurer balance sheets reached \$849 billion in 2024, roughly fourteen per cent of total assets and double the level of a decade earlier. *[FACT — FEDERAL RESERVE BANK OF CHICAGO WORKING PAPER 2025-09]*

Separately managed accounts, hedge funds and pension investors, participating opportunistically through structured vehicles and credit strategies. *[FACT]*

What unites them is the same operational problem seen from the receiving end. Investors receive loan data through "legacy analog workflows including emails, faxes and phone calls," with poor data quality and fragmented technology. *[VENDOR CLAIM]* Buy-side firms find it taxing "to obtain a comprehensive, real-time view of loan details necessary to properly manage their portfolios," lacking timely access to global commitments, outstandings, transaction history, interest accruals, floating-rate details, and notices of principal amortization and prepayments, interest and fee payments, repricings and base-rate contracts. *[VENDOR CLAIM]* Investors "input this new data manually into their own position monitoring systems, sometimes creating errors" that then require reconciliation against the agent's books. *[FACT — REPORTED IN INDEPENDENT FINANCIAL PRESS]*

Two attributable practitioner statements capture the market's own view. A research director at an independent market-structure firm observed that loans "take upwards of 20 days to settle... interesting-looking trades were passed over because settlement was uncertain." *[FACT]* An executive at a financial-messaging firm described manual processes as causing "information mismatch, trade breaks, human errors and lack of appropriate collaboration tools." *[FACT]* The then-head of the principal US trade association cited errors arising from "somebody [who] misunderstood a trading name sent via email or a phone call." *[FACT]*

That last quotation is worth pausing on. The chief executive of the market's standard-setting body describing errors caused by misreading a name in an email is not a vendor pitch. It is the institution responsible for the market's conventions describing what the conventions have not fixed.

The service layer

Between the agent and the investor sits a layer of firms that exists precisely because the handoff between them does not work.

Third-party loan administrators and agency-service providers take on the servicing role for lenders who do not wish to build it, and increasingly for private-credit managers who have no agent bank at all. The largest describes itself as "the largest loan agency service provider on the market," administering \$1.4 trillion of debt assets across loan administration, loan accounting, trade settlement and agency services including payment processing, lender-register maintenance and covenant monitoring. *[VENDOR CLAIM]* Others include corporate-trust arms of major banks, independent agency specialists, and the trustee franchises that serve the CLO market. *[FACT]*

Corporate trustees and CLO administrators occupy a related but distinct position. One provider holds a reported sixty-six per cent trustee share; another approximately thirty-four per cent and the leading agency position. *[VENDOR CLAIM — TRUSTEE SHARE FIGURES COME FROM A COMPETITOR'S RANKING MATERIAL AND SHOULD BE TREATED WITH CAUTION]* A major bank's corporate-trust services business was acquired by a registrar group for approximately \$750 million, completing in November 2021 with roughly 2,000 staff transferred. *[FACT]*


Business-process outsourcers and managed-service providers run the operational work itself, either on the client's platform or their own. *[FACT]*

The strategically important observation about this layer is a division within it, and it will matter in Part Thirteen. Some of these firms are **pure service providers** whose economics improve if better technology exists, because their margin comes from labour they would rather not perform. Others are **platform-owning service providers** whose economics depend on their own software — the administrator that owns a CLO accounting suite, the market-

data firm that sells both a loan-administration platform and its outsourced operation, the incumbent vendor with a middle-office services arm. *[FACT – THE OWNERSHIP FACTS ARE DOCUMENTED; THE STRATEGIC DIVISION IS INFERENCE]* For a new entrant, the first group is a channel and the second is a competitor, and the distinction is not visible from the outside. *[INFERENCE]*

The infrastructure layer

THE FEE, OVER THE LIFE OF THE FACILITY



Above the individual firms sits a thin layer of shared infrastructure, and its ownership history is one of the most important and least examined facts in the market.

12.1 The consolidation under one roof

The market's identifier, pricing, settlement and administration infrastructure passed through a series of acquisitions into a single owner.

Wall Street Office, a family of syndicated-loan portfolio and risk-management products, was developed by a unit of JPMorgan Chase based in Dallas. Markit acquired that unit from JPMorgan in July 2008, completing in September, with a combined footprint cited at more than 400 financial institutions managing over \$1 trillion in assets; the deal value was not disclosed. *[FACT]* ClearPar, the loan-settlement platform, was acquired by Markit in a deal agreed in October 2009 and completed in January 2010 — from Fidelity Information Services, not built in-house. *[FACT]* Markit merged with IHS in an all-stock transaction announced in March 2016 and valued at more than \$13 billion, forming IHS Markit. *[FACT]* IHS Markit merged with S&P Global, completing on 28 February 2022, in an all-stock deal valuing IHS Markit at an enterprise value of approximately \$44 billion. *[FACT]*

The result: loan pricing, loan reference data and identifiers, loan trade settlement, and loan and CLO administration came under one owner. [FACT]

Regulators intervened at the *data* layer and not at the *transactional* layer. The Department of Justice required substantial divestitures and a non-competition waiver to clear the merger, focused on price-reporting-agency overlaps. [FACT] The European Commission cleared the merger conditionally in October 2021 with commitments including divestiture of the leveraged-loan data business; S&P Global consequently sold Leveraged Commentary & Data and its family of leveraged-loan indices — including the principal syndicated-loan index — to a competitor for up to \$650 million, announced in April 2022 and completed mid-year. [FACT]

But ClearPar settlement and Wall Street Office administration were not required to be divested and remain under the same owner. [FACT] That is the crux of the concentration story: antitrust remedies addressed the index and the commentary, and left the transactional plumbing consolidated.

The scale of that plumbing is worth stating. In its 2015 registration statement — a regulated document, not marketing — Markit described ClearPar as the primary platform for electronic confirmation, documentation and settlement of syndicated loans in the United States, settling "substantially all" LSTA leveraged syndicated loans and supporting LMA-market loans. [FACT] Most par and near-par trades settle on ClearPar or on a competing platform. [FACT] Wall Street Office is offered as licensed software and as an outsourced managed service, with vendor claims of hundreds of alternative asset managers and \$4 trillion of assets under administration. [VENDOR CLAIM]

12.2 The notice hub

One figure from this layer deserves separate emphasis because it is the best single measure of the market's manual burden and because it is independently sourced.

Markit's WSOData receives, by email, fax and structured message, more than seven million agent notices on syndicated loans each year, and redistributes over four million of those to clients via electronic straight-through delivery. *[FACT — REPORTED IN THE INDEPENDENT FINANCIAL PRESS]*

That figure dates from 2010. It has not, as far as public sources show, been superseded by a more recent independent count. *[UNKNOWN]* What it establishes is that a single commercial hub sits in the middle of a seven-figure annual flow of unstructured documents whose contents are the market's operational reality — and that the hub's function is to receive, categorise and store them, which is a capture-and-sort layer over an unstructured inflow rather than a fix to the source. *[INFERENCE]*

12.3 The utilities and the networks

Three additional infrastructure positions matter, and each represents a distinct architectural bet. They are examined in detail in Part Six; here they are located on the map.

DTCC Loan/SERV is an industry-owned utility offering daily position and commitment reconciliation between agents and lenders, and a messaging service transmitting industry-standard loan messages in a structured format. *[FACT]* It reached roughly 2,300 funds. *[FACT]* Its messaging portal was acquired by Markit in 2010 and absorbed into WSOData. *[FACT]* The architectural bet: **standardise the messages.**

Versana is a bank-founded platform capturing administrative agents' books-and-records loan data in real time via application programming interfaces and centralising it for lenders. Founded in March 2022 by four banks, live in December 2022. *[FACT]* The architectural bet: **go to the source and bypass the notices entirely.**


AccessFintech, through its Synergy network, is a multi-party data-collaboration layer in which institutions publish transaction, position, cash and reference data into a shared network that normalises it and surfaces the same data and exceptions to all counterparties — explicitly *without*

replacing existing source systems. [FACT] Founded 2016; more than 250 institutions; more than one billion transactions contributed per month; a 2021 integration with the incumbent servicing vendor's lender-data product feeding data back into the core servicing system; and a November 2025 strategic partnership connecting a major buy-side portfolio platform's community to the network. [FACT] Two large banks reported a thirty per cent reduction in trade fails and a seventy-six per cent drop in email traffic after a year on the network. [FACT THAT THE BANKS REPORTED THESE OUTCOMES, RELAYED BY INDEPENDENT TRADE PRESS; THE UNDERLYING MEASUREMENT IS THE PARTICIPANTS' OWN AND HAS NOT BEEN INDEPENDENTLY AUDITED] **The architectural bet: sit alongside everything and reconcile.**

Those three bets — standardise the messages, go to the source, coexist and reconcile — are, across the initiatives examined in this paper, the only three that have been made at scale. Part Six evaluates how each has performed and tabulates them side by side. The reader should carry forward one fact from the third: it is already occupied, by a firm with a bank cap table, an integration with the incumbent servicing vendor, and a decade of network effects. [INFERENCE]

Octaura is the electronic trading venue: announced June 2022, backed by seven major dealer banks and a ratings and analytics firm; the market's first fully electronic syndicated loan trades completed February 2023; loan trading platform launched April 2023; CLO trading live September 2025. [FACT] Its share of secondary loan volume rose from roughly one per cent in the first quarter of 2024 to approximately 4.6 per cent in the first quarter of 2025 and roughly 7.7 per cent by January 2026, with a dealer network that grew from three to around thirty and a buy-side network from thirty-four to more than 170. [VENDOR CLAIM — SELF-REPORTED, DIRECTIONALLY CORROBORATED BY INDEPENDENT TRADE PRESS] Its protocols are request-for-quote and list-based — dealer-intermediated — rather than anonymous all-to-all matching. [FACT]

The consultancy and implementation layer



The incumbent's defensibility is not the software. It is an institutional ecosystem that has accreted around the software over three decades.

A layer that is invisible in most market maps and matters disproportionately, because it is where the incumbent platform is actually installed, operated and kept alive.

The servicing cores of Chapter 12 are not implemented by their vendors alone, and in most institutions they are not principally operated by their vendors either. Around each of them has accreted a services economy — global integrators, platform boutiques, testing firms, managed-service providers, contractor pools and the bank's own long-tenured technology staff — whose collective knowledge of a given installation exceeds the vendor's and cannot be bought from the vendor. This chapter describes that economy as an institutional layer. Chapter 30 then asks what it does to the incumbent's defensibility, and Part Thirteen asks what it does to an entrant's route to market.

13.1 The tiers

TIER	WHO	WHAT THEY SUPPLY	EVIDENCE CLASS
Global systems integrators	Accenture, Capgemini, TCS, Wipro, Cognizant, Infosys, Luxoft/DXC	Programme delivery, target operating model, vendor selection, greenfield implementation, migration, upgrades, testing at scale	Capgemini and TCS are named Leaders in an independent analyst assessment of lending IT services; twenty-eight providers were assessed in that exercise <i>[FACT]</i>
Platform boutiques	LS2 Partners, Adventure Consultancy Solutions, Atyeti	Implementation built on deep single-platform knowledge, advisory, application support and maintenance, staff augmentation	Self-described scope <i>[VENDOR CLAIM]</i> ; the boutique-to-integrator career path is documented <i>[FACT]</i>
Testing and automation vendors	Specialist no-code test-automation firms, generally reached through a boutique partnership	Regression suites for upgrades and migrations — the work that makes a platform change auditable	A named boutique markets automation testing through a partnership with a test-automation vendor <i>[VENDOR CLAIM]</i>
Managed-service and BPO providers	Genpact, Delta Capita, Broadridge, FIS Virtus, Alter Domus, S&P Global WSO Services	Running the operational process on the platform, and in several cases the platform itself, as a service	Self-described scope <i>[VENDOR CLAIM]</i> ; Delta Capita's public role postings include ACBS-to-Loan IQ migration and reconciliation work <i>[FACT]</i>
Independent contractors and alumni networks	Individuals, reached through boutiques and directly	Scarce platform-specific capacity, bought by the day	A boutique markets an explicit "time share" model leveraging platform-knowledgeable independent contractors <i>[VENDOR CLAIM]</i>
The bank's own technology function	Internal	Configuration, interfaces, release management, the accumulated record of why the installation is the way it is	<i>[INFERENCE]</i>

Two documented transactions show the shape of the layer. A global integrator acquired a New York corporate- and commercial-lending consultancy specialising in the dominant platform in **October 2018**, describing the acquired firm's people as having "several hundred years of combined hands-on lending experience"; the acquired firm had been founded in 2013. *[FACT]* A large engineering services firm describes itself as a "**trusted Finastra partner since 2006**" — approaching two decades — across advisory, greenfield implementation, migration, upgrades, managed services and engineering. *[VENDOR CLAIM]* And a specialist boutique founded

after that 2018 acquisition describes itself as "the 'lending' utility company connecting and collaborating with lending ecosystem vendors," offering implementation "built on the complete understanding of the Loan IQ system and the commercial lending business," advisory, automation testing and time-share support. *[VENDOR CLAIM]* Its managing partner previously led the dominant platform's sales and global services delivery at the acquiring integrator, and before that was a managing partner at the firm that was acquired. *[FACT]*

13.2 What the layer actually does

The work is episodic, and that is the reason the capability sits outside the bank rather than inside it.

Selection and target operating model. Before a platform decision, an integrator is engaged to assess the estate and specify the future state. This is the point at which the incumbent's install base is either extended or, rarely, questioned.

Implementation and configuration. The platform is not a product that is switched on. It is a configuration exercise against a specific bank's facility structures, fee conventions, general ledger, entitlements and reporting obligations. The configuration is the asset, and it is built by people who do this repeatedly across institutions.

Migration. Moving from one core to another — most commonly, on the evidence of Chapter 31, from ACBS toward Loan IQ — requires mapping historical positions, accruals and documentation into a new data model and proving the result to audit. The firms that do this work are the only parties with a track record of having done it before. *[INFERENCE]*

Upgrade and regression testing. A servicing core in production is upgraded on the vendor's release cadence and on the bank's own change schedule. Each upgrade requires evidence that nothing broke. Automated regression suites are built and maintained against a specific installation, which makes them an asset of the installation rather than of the platform.

Run and support. Application support and maintenance, defect triage, interface monitoring and month-end assistance are supplied on retainer or by the day, frequently through contractors placed by boutiques.

Capacity. When a bank needs six people who know Loan IQ for nine months, it does not hire six people who know Loan IQ. It rents them. *[INFERENCE]*

13.3 Why the buyer relies on it

Three reasons, and none of them is a preference.

The vendor does not have the people at the necessary scale. A platform vendor sells to hundreds of institutions and cannot staff each installation's configuration, integration and support with its own employees. The partner network exists because the delivery model requires it. *[INFERENCE]*

The bank does not have them either, and increasingly cannot keep them. Chapter 66's evidence on loan operations — one to two years to productivity against twenty to seventy per cent attrition *[VENDOR CLAIM]* — describes the operating side of the same problem on the technology side. Platform-specific knowledge takes years to build and leaves at market rates.

Migration is the one activity that most requires the knowledge nobody has in reserve. A bank contemplating a core change needs, precisely at that moment, more platform expertise than it has ever needed in steady state — and it needs it in both the outgoing and the incoming system simultaneously. The market's answer is to buy it, and the firms selling it are the firms whose revenue depends on the incumbent estate. *[INFERENCE]*

13.4 What the layer does to the incumbent's position

It cuts two ways, and both are consequential.

It lowers the incumbent's operating risk. A platform surrounded by a deep external services market is easier to run, easier to staff, easier to upgrade and easier to recover when something goes wrong. Scarcity of expertise

would be a reason to leave a platform; abundance of purchasable expertise removes that reason. *[INFERENCE]*

It raises the cost and the perceived risk of replacement. The same abundance is asymmetric. There is a mature market for people who can implement, migrate to, test and support the incumbent platforms. There is no equivalent market for any challenger, because a challenger with few installations cannot have generated one. A bank replacing its core therefore trades a liquid expertise market for an illiquid one, and does so at the exact moment it most needs depth. *[INFERENCE]*

And the ecosystem's own economics run one way. These firms bill hours against Loan IQ and ACBS. Their revenue is a function of the incumbent's persistence, not of its quality. *[INFERENCE]*

13.5 The stronger proposition, and its limits

The proposition worth stating plainly is this:

The incumbent's defensibility is not the software. It is an institutional ecosystem — software, configuration, internal knowledge, external expertise, implementation firms, operating procedures, contractors, testing infrastructure and career capital — that has accreted around the software over three decades. *[INFERENCE]*

Chapter 30 formalises the part of that ecosystem which sits inside the bank, as Layers 2 through 4 of the moat. This chapter's addition is that a substantial part of it sits *outside* the bank, is purchasable, and is therefore usually left out of the switching-cost arithmetic — where it belongs on the cost side, because the buyer has to re-acquire it in a market that does not yet exist for the alternative. *[INFERENCE]*

Three limits keep the proposition honest.

The ecosystem is not a decision-maker. These firms advise, implement and operate; they do not hold a veto over a bank's platform strategy, and no public evidence shows one attempting to block a migration. The

mechanism is economic gravity, not obstruction. *[INFERENCE]*

Its influence is not measured. No public source quantifies how much of a bank's servicing-platform total cost of ownership is external services, or how a migration's budget divides between licence, integrator and internal cost. That figure would materially sharpen Chapter 31, and it does not exist publicly. *[UNKNOWN]*

And most of the evidence here is self-described. Firm capability claims are marketing. What is independently documented is narrower: the 2018 acquisition and its stated rationale; an analyst assessment naming twenty-eight providers in lending IT services and several as Leaders; a boutique-to-integrator-to-boutique career path; and role postings for migration and reconciliation work between the two incumbent cores. *[FACT]* The structural argument rests on that narrower set. *[INFERENCE]*

13.6 Coexistence and replacement are different businesses for this layer

The distinction matters more than it first appears, and Part Thirteen turns on it.

	A PRODUCT THAT COEXISTS WITH THE INCUMBENT	A PRODUCT THAT REPLACES THE INCUMBENT
Global integrator	New implementation and integration scope; extends the programme	Competes with its own modernization narrative; large but risky programme
Platform boutique	Something new to implement, test and support on top of an existing base	Destroys the annuity outright
Testing and automation vendor	More surface to regression-test	Rebuild from zero, against an unfamiliar platform
Managed-service provider without its own platform	A better middle office to operate	Neutral to positive
Managed-service provider with its own platform	Direct competitor	Direct competitor
Contractor pool	Additional demand for existing skills	Demand for skills it does not have

[INFERENCE]

The pattern is that additive technology has a natural distribution channel in this layer and replacement technology has a natural adversary in the same layer — with one exception that should not be overlooked, since the managed-service firms that own their own servicing technology are competitors in either case. *[INFERENCE]* The boutiques already resell third-party technology into their installed base, which is a working template rather than a hypothesis. *[FACT]*

Two threads run forward from here. Chapter 34 takes up what it means that the deepest platform knowledge now lives outside the platform's vendor. Chapter 95 takes up whether this layer is a viable route to market for a firm that has no institutional relationships of its own — and finds that it is the best available answer to the anchor problem, at the price of margin and direct customer contact.

Who pays whom

The market's fee architecture is mostly private, and where it is visible it is instructive.

14.1 The fee stack

FEE	WHAT IT COMPENSATES	WHAT IS PUBLICLY DOCUMENTED
Arrangement / underwriting	Structuring, underwriting, syndicating	Leveraged loans commonly 1–5% of commitment; higher for acquisition and restructuring financings, lower for seasoned refinancings; investment-grade facilities pay little <i>[INFERENCE FROM PRACTITIONER CONVENTIONS, NOT AUDITED DATA]</i>
Upfront fee / original issue discount	Lender commitment at close	Quoted in points of par; treated as OID for tax <i>[FACT]</i>
Agency fee	Running the loan after close	Flat annual dollar amount, set in a private fee letter, no public range <i>[FACT]</i>
Commitment fee	Undrawn portion of a revolver	Approximately 0.25–1.00% per annum <i>[FACT]</i>
Facility fee	Entire committed amount, common in investment-grade revolvers	Order of 10–18 basis points <i>[FACT]</i>
Ticking fee	Gap between commitment and funding	Step-up structure, commonly 0% for days 0–30, 50% of margin for 31–60, 100% thereafter; accrual start historically 90–180 days post-allocation, compressing toward 30 or 15 <i>[FACT]</i>
Amendment / consent fee	Consenting lenders' agreement to change terms	Percentage of consenting exposure; deal-specific <i>[FACT]</i>
Assignment fee	Agent processing of a lender-to-lender trade	\$3,500 standard form; some agreements \$4,500; waivable at agent's sole discretion <i>[FACT]</i>

The peer-reviewed anchor for the whole structure is a 2016 *Journal of Finance* study establishing that more than eighty per cent of US syndicated loans carry at least one fee type, that contracts specify a menu of spreads and fees rather than a single price, that fees can exceed interest payments for some loans, and that the standard all-in-drawn-spread measure

understates the true cost of borrowing. *[FACT — PEER-REVIEWED; NOTE THAT THE SPECIFIC ">80% / FEES CAN EXCEED INTEREST" FIGURES SIT BEHIND A PAYWALL AND SHOULD BE ATTRIBUTED TO THE PAPER RATHER THAN QUOTED AS INDEPENDENTLY RE-VERIFIED]*

That last point is the reason the fee structure belongs in a paper about operations rather than one about pricing. **A loan is not priced by a rate. It is priced by a menu, and the menu's items are events that must be calculated, notified, collected and distributed by the agent's operations team over the life of the facility.** *[INFERENCE]* Every fee in the table above is a servicing event. The market's pricing complexity and its operational complexity are the same phenomenon viewed from two sides.

14.2 Concentration

Arrangement fees flow disproportionately to a small number of institutions.

In 2024 US leveraged loans, Bank of America ranked first with 953 deals and more than \$172 billion of credited bookrunner volume; JPMorgan second with 961 deals and roughly \$167 billion; Wells Fargo third at roughly \$137 billion; Goldman Sachs fourth at roughly \$111 billion. The top four accounted for nearly thirty per cent of total bookrunner volume — in a year when the three leaders each *lost* share. *[FACT]* In all-in global syndicated loans including investment grade, JPMorgan tops the volume rankings. *[FACT]*

The correct framing is therefore "a top-two duopoly that alternates first and second depending on which league table," not a fixed ordering. *[FACT]* A commonly repeated claim that JPMorgan has held the number-one arranger position for twenty-six consecutive years appears on the bank's own materials citing a data provider's league tables; the rank itself is well supported, the "26 years" phrasing is the bank's own. *[VENDOR CLAIM]*

Global syndicated lending proceeds rose thirty-two per cent to \$5.9 trillion in 2024, with the United States at \$3.6 trillion, Europe at \$1.0 trillion and Asia-Pacific at \$567.4 billion. *[FACT]* A frequently circulated figure that

JPMorgan generated approximately \$1.8 billion of fees on \$433.6 billion of 2024 deals, within an industry total of roughly \$29 billion across 588 lead arrangers, recurs consistently across secondary reporting but could not be traced to a primary publisher; **this paper does not rely on it.** [UNKNOWN]

14.3 What concentration does to price

There is peer-reviewed evidence that arranger concentration raises borrower cost. A 2020 study in the *Review of Quantitative Finance and Accounting* finds that syndicated loan prices — fees, spreads and rates — are positively associated with the concentration of both the borrower's and the lead arranger's markets, that pricing is more sensitive to lead-arranger concentration than to borrower-market concentration, and that a rise in lead-arranger market concentration raises the spread on a typical loan by approximately 4.15 basis points, equivalent to roughly \$152,000 of additional cost for a typical corporate borrower. [FACT]

The study's most interesting finding is a nuance usually omitted when it is cited: **the relationship reverses when an investment bank or non-bank leads the syndicate.** [FACT] The concentration premium is a commercial-bank-oligopoly phenomenon, not a general property of concentrated arranging. That qualification matters for Part Eight, because the growth of non-bank arranging in private credit is, on this evidence, a competitive rather than an anticompetitive development. [INFERENCE]

A map of authority

THE AUTHORITY MAP

LAYER	ROLE	AUTHORITY · CROSS-INSTITUTION NEED
1 · Borrower	Authoritative on its own condition and on nothing else.	
2 · Agent	Authoritative where the agreement makes it authoritative — most importantly on lender-of-record state; not authoritative for every economically relevant fact.	
3 · Lenders	Authoritative on their own books, derivative on everything else.	
4 · Administrators, custodians, trustees	Records for accounting, custody, valuation, compliance, waterfalls and reporting — some derivative, some carrying their own authority; not simply inferior copies.	
5 · Market infrastructure	Authoritative on nothing about any individual loan and indispensable to the market's functioning.	

Bar lengths are illustrative of direction, not measured quantities.

The party whose need for a cross-institution view is smallest holds the authoritative data; the parties whose need is greatest hold derivative copies whose accuracy they cannot verify without asking the party that is not compensated for making verification easy.

Pulling the anatomy together, the market can be described as five layers with a specific and unusual property: **authority decreases as need increases.**

Layer 1 — The borrower originates the underlying facts: the drawdown request, the compliance certificate, the financial statements, the request for a waiver. It is authoritative on its own condition and on nothing else.

Layer 2 — The agent converts those facts into the portions of loan state the credit agreement assigns to it: the register, specified interest and fee calculations, payments, assignments, votes and the notices through which those determinations are communicated. It is authoritative where the agreement makes it authoritative — most importantly on lender-of-record state — but it is not authoritative for every fact economically relevant to the loan.

Layer 3 — The lenders hold economic interests in the loan and maintain their own records of positions, cash and accruals. They are authoritative on their own books and derivative on everything else. When their records disagree with the agent's register, their records do not govern the loan's operative lender-of-record state. They may ultimately prove the agent's record wrong, but the burden is on them to challenge and correct the authoritative record.

Layer 4 — The administrators, custodians, fund accountants and trustees maintain records for accounting, custody, valuation, compliance, waterfalls and reporting. Some are derivative of agent and lender data; others answer questions for which those functions possess their own responsibility or authority. They are not simply inferior copies of the agent's record.

Layer 5 — The market infrastructure — pricing services, identifier registries, settlement platforms, data utilities — assembles cross-institution views from the layers below. It is authoritative on nothing about any individual loan and indispensable to the market's functioning.

The pathology is now visible in a single sentence. **The party whose need for a cross-institution view is smallest holds the authoritative data; the parties whose need is greatest hold derivative copies whose accuracy they cannot verify without asking the party that is not compensated for making verification easy.** *[INFERENCE]* The agent of course needs its own consolidated record and has one. What it does not need, and therefore has no reason to invest in, is everyone else's ability to reproduce it.

Part Three shows what that costs, event by event. Part Ten asks what could be done about it. But the reader should carry this map through everything that follows. Many of the market's operating problems arise at the boundary between an authoritative source and the institutions that must reconstruct, interpret or act on what it states. Later parts show why the modernization failures cannot all be reduced to that boundary: some failed on coordination, some on incentives, some on legal authority, and some because they solved a different problem from the one the market actually had.

III

PART THREE

Mechanism

How the loan lifecycle actually works

This part is the technical heart of the paper. It describes, event by event, what actually happens to a syndicated loan from mandate to maturity, who does it, what system it touches, and where it breaks.

The description is deliberately granular, because the recurring failure of outsiders analysing this market is to reason about it at the level of abstraction at which it is usually discussed. "Loans are manual" is true and useless. What is useful is knowing that a lender's interest payment under a daily-compounded overnight reference rate cannot be calculated until roughly four business days before it is due, that the calculation changes if the lender's position changed mid-period, and that the resulting notice goes out to a syndicate that may contain several hundred entities. That is the level at which one can distinguish a problem software can solve from a problem it cannot.

The part closes with Chapter 26, which decomposes the market's complexity into five categories and argues that they have entirely different susceptibility to automation. That chapter is the analytical pivot of the paper.

Origination and primary syndication

The documented sequence from mandate to close runs: mandate and term sheet, commitment letter, information memorandum, lender invitations and confidentiality agreements, bank meeting, draft facility documents, commitment letters returned, allocations, signing and closing, drawdown. From launch to close the process typically takes **ten to twelve weeks**. *[FACT]*

Unpacked into its manual steps:

1. Arranger and borrower negotiate and execute a term sheet and commitment letter setting pricing and structure.
2. The arranger prepares an information memorandum and financial materials and stands up a secured deal site.
3. Lender invitations go out; signed confidentiality agreements are collected before data-room access is granted.
4. A bank meeting is held; lender questions are fielded and answered.
5. Draft facility documents are distributed; commitment letters are collected from interested lenders.
6. Final commitments are allocated. If the book is oversubscribed the arranger prorates allocations — a manual judgment and modelling step, not a mechanical one.
7. Bank allocation, funding and closing memoranda are prepared and distributed.
8. Facility documents are executed at close; the initial drawdown is processed.

Two aspects of this sequence deserve emphasis.

Per-lender onboarding is a repeated manual step. The agent coordinates loan assignments including lender onboarding and record-keeping, and know-your-customer diligence is performed per syndicate member. *[FACT]* There is no public figure for the cycle time or headcount cost of per-lender

KYC; practitioner sources describe it as a settlement drag without quantifying it. *[UNKNOWN]* This gap recurs throughout Part Three and is worth noting as a class: the market's operational burdens are consistently described qualitatively and almost never measured publicly, which is itself a finding.

Allocation is judgment, not calculation. When a book is oversubscribed, the arranger decides who gets what, weighing relationship, future business, the composition of the eventual lender group, and the borrower's preferences. No system automates this and none should. *[INFERENCE]* It is the first of several points in the lifecycle where the work is irreducibly discretionary.

The ten-to-twelve-week timeline is the baseline against which private credit's principal competitive claim — speed and certainty of execution — should be read. A direct lender that can commit in two weeks is not offering a marginally faster syndication; it is offering a different product. *[INFERENCE]*

Closing, funding, and the birth of the register

At close, the facility comes into existence as a legal object and as an operational one.

The agent establishes the register: the list of lenders, their commitments, and their principal amounts, to be updated from each assignment and acceptance thereafter. *[FACT]* Every subsequent event in the loan's life — every payment, every vote, every notice — is computed against this record.

What the register holds. At its core it is a small set of facts: which legal entities are lenders, in what capacity, for what commitment on each tranche, with what principal currently outstanding, and — for the purpose of paying them — where the money goes. It is maintained by the agent and updated from each assignment and acceptance. *[FACT]* It is not a copy of the credit agreement and it is not the loan's economics; it is the answer to one question, *who owns what*, and it is the answer to that question that everything else is computed against.

What else is fixed at close, and is not in the register. The executed agreement itself, with the negotiated terms Chapter 36 describes; the margin grid and fee schedule the agent will apply; the drawdown and funding mechanics; the notice list — who receives what, by what channel; and the consent regime that will govern every subsequent transfer (Chapter 22.3). These live in a document, not in a data structure, and the operational problem of Part Three is in large part the problem of getting them *out* of the document and into systems, correctly, once per institution. *[INFERENCE]*

And this is the moment the parallel representations begin. From the instant the register exists, other institutions begin maintaining the representations they need: a lender's position, a custodian's cash record, a

fund administrator's books, a CLO trustee's waterfall inputs, a pricing vendor's file and a rating agency's model. Many depend on information originating with the agent; they are not all copies of the same thing.

It is worth stating plainly what is singular and what is not. **From the moment of close, lender-of-record state has one authoritative record: the register maintained by the agent. But the loan as an economic and institutional object immediately exists in multiple representations, because different participants have different functions and answer different questions about it.** *[INFERENCE]*

The consequential part is what is *not* created alongside the register. Nothing at close establishes a comprehensive mechanism by which every downstream representation remains synchronized with the authoritative facts on which it depends. For the agent-controlled facts, the principal propagation mechanism is a stream of notices and other communications. *[FACT]* Every downstream participant must therefore update its own state from events communicated by another institution.

This is the origin of much of the reconciliation burden in Chapter 25 and the settlement friction in Chapter 22. The problem is not that the market has no authority. It is that authority and information are distributed, while propagation between them remains incomplete, delayed and semantically difficult. *[INFERENCE]*

The notices are the subject of the next chapter.

Servicing I: the notice as the unit of work

1. Borrowing / drawdown — pro rata funding by a specified date and time.	2. Interest-rate-set — the reference rate plus margin, per period.
3. Interest / fee payment — advance notice and remittance.	4. Rollover / continuation / conversion — irrevocable, effective on notification.
5. Prepayment / paydown — the control record cash is reconciled against.	6. Commitment reduction / amortization
7. Payment-in-kind — interest capitalised to principal.	8. Amendment / waiver / consent — requests, results.
9. Assignment / position transfer — the classic position break.	10. Default / acceleration / enforcement
11. Erroneous-payment — conclusive absent manifest error.	

If one wants a single unit of account for the operational burden of this market, it is the notice.

18.1 The taxonomy

The notice set resolves into roughly eleven types, each anchored in a credit agreement's defined terms. *[FACT]*

1. **Borrowing or drawdown notice.** The borrower requests a borrowing; the agent notifies the syndicate; each lender is generally obliged to make its pro rata share available to the agent by a specified date and time. *[FACT]*
Appears in credit agreements as a Request for Credit Extension, Notice of Advance, Notice of Borrowing or Committed Loan Notice.
1. **Interest-rate-set notice.** On each interest period the agent sets and communicates the reference rate plus margin. *[FACT]*

1. **Interest and fee payment notices.** Advance notice of interest, commitment, facility, utilisation and letter-of-credit fees due, and of distributions remitted. *[FACT]* The standardisation literature notes that "the term facility fee is used by banks to denote either commitment fee or utilisation fee" — which is exactly why these notices resist automation. *[FACT]*
1. **Rollover, continuation or conversion notice.** At the end of an interest period a loan is continued into a new period or converted between rate options, via a Notice of Continuation or Notice of Conversion, each irrevocable and effective on notification. *[FACT]*
1. **Prepayment or paydown notice.** Voluntary or mandatory prepayments under a defined Notice of Prepayment. *[FACT]* The post-Revlon erroneous-payment framework treats a payment that "does not match the amount set forth in a prepayment notice" as presumptively made in error — establishing the prepayment notice as the control record against which cash is reconciled. *[FACT]*
1. **Commitment-reduction and amortization notice.** Scheduled amortization, commitment reductions and shifting repayment schedules. *[FACT]*
1. **Payment-in-kind notice.** Interest capitalised to principal rather than paid in cash, communicated as a servicing and position-change event. *[FACT — TREATED IN THE LITERATURE AS A SUBTYPE OF THE POSITION-CHANGE FAMILY RATHER THAN AN INDEPENDENTLY DOCUMENTED NOTICE TYPE]*
1. **Amendment, waiver and consent solicitation and results.** The agent distributes requests, collects consents and communicates results. *[FACT]*
1. **Assignment and position-transfer notice.** When a lender sells, the agent processes the assignment. Misalignment here is the classic position break. *[FACT]*
1. **Default, acceleration and enforcement notices.** *[FACT]*
1. **Erroneous-payment notices.** Post-Revlon model language provides that "a notice of an erroneous payment from an administrative agent to a lender is conclusive absent manifest error," with a presumption that off-schedule payments must be returned. *[FACT]*

18.2 The volume and the medium

The scale figure has already been given and bears repeating in context: a single commercial hub receives more than seven million agent notices on syndicated loans each year by email, fax and structured message, and redistributes more than four million electronically. *[FACT]*

The medium is the problem. "The industry continues to rely on faxes as its primary communication method" — a 2008 characterisation from an industry utility's white paper. *[FACT]* Fourteen years later, the launch coverage of a bank-founded data platform described the market as moving "closer to ditching faxes." *[FACT]* Investors "get updates on their positions through a hodgepodge of emails, instant messages, phone calls, and even the occasional fax." *[FACT]*

18.3 Why the notices resist standardisation

This is the crux, and it is more interesting than "the industry is old-fashioned."

The notices are "essentially unstructured and unorganized," "non-standardized in terms of business language and content," and "recorded both by agent bank and participant lenders independently." *[FACT]* But the deeper problem is **semantic, not syntactic**. Because there is no enforced standard, "one agent bank may choose to provide start date and end date, while another bank may provide start date and tenor," and "facility fee" means commitment fee to one sender and utilisation fee to another. *[FACT]*

That distinction is doing a great deal of work and is worth making explicit.

A notice standardizes transmission. It does not standardize interpretation.

A messaging standard can require everyone to use the same envelope, fields or syntax. It cannot, by itself, make the economic meaning of those fields identical across bespoke agreements. "Facility fee" can still refer to different economics. A date can still have meaning only when read against

a particular provision. A notice can therefore be perfectly transmitted and still require contextual interpretation before a receiving institution knows what state should change. *[INFERENCE]*

This is why the messaging-standard project in Chapter 38 could improve syntax without eliminating semantics; why the identifier project in Chapter 37 produced cross-reference rather than universal identity; and why the extraction problem in Part Nine is fundamentally a problem of meaning rather than optical recognition.

18.4 Why the notice is the choke point

Three linked facts establish it.

The notice flow is the loan's operational data. Every position, cash movement and rate is transmitted as a notice; when notices lag or mismatch, records diverge. The diagnosis from a platform built specifically to address this is that "the majority of cash flow mismatches from agents to lenders result from position discrepancies," identified through "display of each digital transaction notice causing such differences" — the notice is the root-cause record for a break. *[VENDOR CLAIM]*

Breaks originate from unsynchronised notices. The canonical failure modes are position breaks — position changes and other vital information not synchronised between parties — and cash breaks from payment delays, both traced to manual, fax-based notice flow. *[FACT]* The post-Revlon erroneous-payment regime is, in substance, a legal backstop for when a payment does not match the notice. *[FACT]*

Notices originate from the agent's core servicing system. The agent generates them from its loan servicing platform, and the dominant platform's modules explicitly cover bills, notices and enhanced notice generation. *[VENDOR CLAIM]*

Therefore: **any modernization effort must capture data at or from that origin.** *[INFERENCE]* There are only two ways to do it — go to the source system via interfaces, or normalise the notices after they leave it — and the market has tried both. That is the architectural fork of Chapter 44 and Chapter 46.

FIGURE 2

One notice, many reconstructions

The envelope can be standardized. The meaning still has to be reconstructed. One event generates a notice to every affected party; each recipient interprets the same notice through its own systems, data, policies and obligations.

1 · SOURCE EVENT

Credit agreement Legal terms	Borrower action Draw, prepayment	Agent calculation Amounts, dates, allocations	Market condition Rate reset, benchmark	Other events Amendment, consent, default
--	--	---	--	--



2 · NOTICE GENERATION — AGENT / LOAN SERVICING

Notice templates · data pulls (loans, positions, parties) · business rules and calculations · validation and approval · distribution lists · multi-channel delivery.



3 · STANDARDIZED ENVELOPE — COMMON FORMATS, MULTIPLE CHANNELS

PDF · email attachments · portal (web / vendor) · structured / API delivery · fax.

Notice families discussed in Chapter 18. Interest payment · principal payment · rate reset · borrowing or draw · prepayment · fee notice · amendment · consent request · waiver · assignment · covenant delivery · default or reservation of rights · maturity or extension · tax and regulatory notice.



4 · RECIPIENTS INTERPRET AND PROCESS — THE SAME NOTICE

Lender Update position, accrue interest, check covenants → position record	Custodian Confirm cash receipt, update asset records → custody record	Fund administrator Allocate to funds, calculate fees, update NAV → fund records	CLO trustee Test covenants, track triggers → compliance record
Fund accountant Book interest, accrue income, apply accounting policy → accounting record	Risk / analytics Update exposure models, limits and stress tests → risk record	Regulator Map to the reporting framework → regulatory record	Investor Include in investor reporting and performance → investor record

THE SAME FIELDS. DIFFERENT QUESTIONS.

Parsed directly: payment amount · effective date (usually deterministic).

May require interpretation: applicable rate (term interpretation) · day-count convention (contract-specific) · allocation (position-specific).

Requires judgment or authority: fee treatment (accounting and tax policy) · consent threshold (legal interpretation) · who is entitled (register and timing dependent) · whether the event is valid (evidence and contract).

Standardizing the message does not standardize the meaning. Every recipient reconstructs the event inside its own systems, policies and obligations.

A single economic event produces a notice that every recipient must reconstruct inside its own systems, policies and obligations. The envelope can be standardized; the meaning still has to be interpreted, which is the distinction drawn in §18.3. Notice families shown are those discussed in Chapter 18. [FACT]

Servicing II: interest, resets, fees, and the awkward arithmetic of SOFR

The servicing workstreams, enumerated:

1. **Interest and fee calculation and notices**, per interest period, per facility, across the lender group, with notices and invoices issued.
2. **Rate resets**, materially harder since the reference-rate transition.
3. **Rollovers and utilisations**, processing drawdown and rollover requests and interest-period elections.
4. **Amendment, waiver and consent processing**, including vote tabulation.
5. **Position reconciliation across lenders**.
6. **Wire transfers and cash reconciliation**.
7. **Notices to the syndicate**, transmitted "sometimes to several hundred investors" on amortizations and covenant amendments. *[FACT]*
8. **Borrowing-base updates and covenant monitoring**, with borrowers submitting compliance certificates and agent and lenders checking calculations.

Point 2 deserves elaboration because it is the clearest example of an external change that made servicing harder rather than easier.

19.1 The SOFR problem

Under LIBOR, a term rate was fixed at the start of an interest period. Interest was therefore knowable in advance.

Most new US syndicated loans now reference forward-looking Term SOFR, which, like LIBOR, is fixed at the start of the interest period. *[FACT]* Where a loan instead uses daily SOFR compounded in arrears, as the fallback provisions contemplate, "the full period's interest rate is not known until the end of the interest period." *[FACT]* Such loans use a lookback of roughly five business days, so interest is communicated only about four business

days before payment. [FACT] Additional complexity arises "when a lender's position changes within the period." [FACT] From the investor's side: they "can't be sure exactly what interest payment to expect, based off overnight rates accruing over the preceding 90 days." [FACT]

Consider what that combination means operationally. The agent must calculate a compounded rate over a period, apply it to positions that may have changed during the period, produce per-lender amounts, notify several hundred entities, and do it inside a four-business-day window. Every lender must then independently verify the amount against its own position record — which may itself be wrong, because it depends on assignments that may not have settled. And the whole sequence repeats every interest period, on every facility.

This is the market's hardest recurring calculation, it is performed under time pressure, it is performed on data whose accuracy the recipient cannot independently confirm, and it is performed by a function that has been relocated toward lower-cost labour markets. [INFERENCE] It is not surprising that it produces breaks. It would be surprising if it did not.

Amendments, waivers, consents: the vote as an operational object

An amendment is where the loan's contractual nature becomes an operational burden.

Obtaining majority lender consent is "an administratively burdensome task when a syndicated lender group can comprise many hundreds of entities," typically requiring at least 50.1 per cent for ordinary amendments and higher thresholds for sacred rights. *[FACT]* Some deals have 100 to 300 lenders. *[VENDOR CLAIM]*

The operational reality: the agent must identify who the lenders currently are — which depends on the register, which depends on assignments having been processed — distribute the solicitation, track responses, chase non-responders, compute the vote against the correct thresholds and the correct denominators, and communicate the result. If a position transferred mid-solicitation, the agent must determine which party holds the vote, and the trading documentation contains machinery for passing amendments and waivers arising between trade date and settlement date through to the buyer. *[FACT]*

The reason to dwell on this is that it exposes a property of the market that pure data-plumbing analyses miss. **A vote is not a data record; it is a legal act with a deadline, performed by identified parties, whose validity depends on the register being correct at a moment in time.** *[INFERENCE]* It cannot be reconstructed after the fact from a data feed. Any system that aspires to authority over loan state must be able to handle the vote, and the vote is where contractual and operational reality are least separable.

Secondary trading: price discovery without a tape

The secondary market's structure explains a great deal about why the rest of the market looks as it does.

21.1 Size and shape

Annual secondary loan trading volume reached a record of approximately **\$971 billion in 2025**, eighteen per cent above the previous record of \$824 billion in 2022 and above 2024's \$821 billion. Trailing-twelve-month volume surpassed \$1 trillion for the first time in the first quarter of 2026. *[FACT]*

Within 2025, the second quarter was a record at roughly \$262 billion; the third quarter then fell about eighteen per cent — the market is volatility-driven rather than steadily growing. *[FACT]*

The structure is dealer-intermediated over-the-counter. A small set of bank trading desks make markets; the buy-side counterparties are CLOs, loan mutual funds and ETFs, hedge funds and separately managed accounts.

[FACT] The market splits into par/near-par and distressed segments, each with its own documentation and settlement conventions. The precise 2025 par/distressed mix sits inside member-gated trade-association reports and is not publicly available. *[UNKNOWN]*

21.2 How a trade happens

Price discovery. There is no consolidated tape. Discovery is dealer-driven: dealers circulate "runs" of indicative bids and offers, and buy-side sellers run bids-wanted-in-competition — a timed auction where a list of names goes to dealers, price talk forms from where comparable paper last traded, and the high bidder wins. *[FACT]*

Execution. Trades are still largely agreed bilaterally by phone or instant message, in what independent trade press describes as "a highly manual, labor-intensive process." *[FACT]*

Settlement method. Either assignment — the buyer becomes lender of record on the agent's books — or participation, where the seller remains lender of record and the buyer takes a contractual interest, retaining seller credit exposure. *[FACT]* Delayed-compensation rules generally do not apply where the parties agreed at trade time to settle via participation. *[FACT]*

21.3 The transparency gap

Loans are not securities under US law and are not reported to the corporate-bond transaction-reporting facility, so there is no post-trade transaction feed comparable to corporate bonds. *[FACT]* Dealer markets in leveraged loans have been characterised in academic and legal analysis as "opaque, which impairs price discovery." *[FACT]*

The market marks instead from vendor pricing derived from dealer quotes: one major provider covers more than 8,000 leveraged loan facilities sourced from sixty-plus sell-side desks; two competitors sell equivalent evaluated pricing. *[VENDOR CLAIM]* These are quote-derived, not trade-derived. *[FACT]*

Who benefits from opacity is not a rhetorical question; it has an empirical answer. Academic work on the introduction of post-trade transparency in corporate bonds found "substantial reductions in the bid-ask spreads" investors pay. *[FACT]* The counter-argument in the literature is that transparency makes dealers reluctant to carry inventory. *[FACT]* Applied to loans: because pricing flows from dealer quotes rather than a public tape, dealers hold the information advantage on where paper is really trading, and the buy side pays wider, less observable spreads. *[INFERENCE]*

Chapter 53 develops the consequence. If transparency transfers value from dealers to investors, and dealers control the market's infrastructure decisions, then the market's persistent opacity is not an oversight. *[INFERENCE]*

21.4 Electronification

Loan e-trading share rose from roughly one per cent of secondary volume in the first quarter of 2024 to approximately 4.6 per cent in the first quarter of 2025 and roughly 7.7 per cent by January 2026, with quarterly volume of \$19.6 billion in the first quarter of 2026 and a March 2026 record of \$7.7 billion across 1,322 loans. *[VENDOR CLAIM — SELF-REPORTED BY THE VENUE,*

DIRECTIONALLY CORROBORATED IN TRADE PRESS]

Two structural observations.

The protocols are dealer-intermediated. The venue's model is request-for-quote and list-based: the buy side pings multiple dealers and transacts with the chosen dealer. It electronifies the workflow while keeping the dealer in the middle; it is not anonymous all-to-all matching. *[FACT]* A frequently repeated claim that "only about one per cent of the market trades via an electronic matching engine" cannot be cleanly verified as stated; what is verifiable is that total electronic penetration is single-digit and that the dominant venue's protocols are not central-limit-order-book matching. *[UNKNOWN]*

The benchmark shows how far this has to run. Roughly forty-three per cent of US investment-grade and high-yield corporate bond volume traded electronically in 2024 — forty-eight per cent investment grade, thirty-three per cent high yield — up from forty-one per cent in 2023. *[FACT]* Loan electronification is therefore at roughly one-sixth to one-tenth of the corporate bond level. *[INFERENCE]*

The instructive part of the bond comparison is the causal sequence. Bond electronification progressed alongside the arrival of a public tape and protocol innovation. *[FACT]* Loans have the protocol innovation and not the tape. *[INFERENCE]* Whether the tape is a precondition or merely a correlate is the kind of question this paper flags rather than settles. *[UNKNOWN]*

Settlement: T+7, T+20, and what actually happens

SETTLEMENT, BY ASSET CLASS



Bar lengths scaled to days; a settlement convention is a target, not an outcome — the loan average is realised time, the others are standards.

The market's settlement conventions and its settlement reality are two different things, and the gap between them is the market's most legible failure.

22.1 The conventions

The standard is **T+7 business days for par and near-par trades and T+20 for distressed**. *[FACT]* Settlement beyond those windows triggers delayed compensation — an interest and fee true-up putting buyer and seller in the economic position they would have held had the trade settled on time. *[FACT]*

Since 1 September 2016 the regime has been "requirements-based" rather than "no-fault": the buyer must satisfy basic requirements by roughly T+5 — execute documents, be able to settle — to earn accruals from T+7. *[FACT]* The change was made because the trade association concluded the previous no-fault regime was counterproductive: "Buyers may not be motivated to settle trades when they are guaranteed delayed compensation." *[FACT]* Buy-in/sell-out provisions let a performing party terminate a stalled trade and cover, but only at T+20 — "too long after a trade date to mitigate delayed settlement risks in less dramatic situations." *[FACT]*

22.2 The reality

Average par settlement ran approximately **19.5 days in 2023**, a seven per cent improvement on 2021–22. *[FACT]* The best-documented distribution figures, from 2021 data, are that **only about twenty-nine per cent of par trades settled within T+7 and about twenty-seven per cent settled later than T+20**. *[FACT — THIS SPECIFIC PAIRING IS CITED IN LAW-FIRM ANALYSIS; IT COULD NOT BE TIED TO A SPECIFIC DATED TRADE-ASSOCIATION PUBLICATION AND ONE RESEARCH PASS FLAGGED IT AS UNCONFIRMED, SO IT SHOULD BE ATTRIBUTED TO THE LAW-FIRM SOURCE RATHER THAN PRESENTED AS A TRADE-BODY STATISTIC]* A separate vendor factsheet cites "T+16 average settlement (par)." *[VENDOR CLAIM]* A 2016 trade-association study found only twenty-three per cent of par trades settling within T+7. *[FACT]*

The two metrics — a T+16 or T+19.5 *average* and a T+20 *distressed standard* — are different things and are frequently conflated in casual usage. *[FACT]* Current 2025 mean and median settlement days are inside member-gated studies and are not public. *[UNKNOWN]*

The comparison that makes the point: **US equities settle T+1; corporate bonds settle T+1 or T+2; syndicated loans average roughly three weeks**. *[FACT]*

22.3 Why

The friction points are documented specifically, and none of them is a technology gap.

- Electronic settlement platforms "may lack automatic feeds from broker systems or have credits not built into the system." *[FACT]*
- "Allocations must often be provided or entered into an electronic settlement platform manually" and "may not be available on the trade date." *[FACT]*
- Confirmation disputes and slow internal closer and signatory review. *[FACT]*
- "Some institutions will not sign an Assignment Agreement until the Trade Confirmation has been fully executed." *[FACT]*

- "Administrative agents often do not have the capacity to review every trade within one business day." *[FACT]*
- Intentional postponement for cash-flow and balance-sheet management. *[FACT]*
- Borrower and agent consent requirements, disqualified-lender list checks, and per-lender KYC. *[FACT]*
- Consent denials and slow consent responses, a recognised and rising drag. *[FACT]*

Read that list again and note what it contains: a consent regime, a documentation regime, an approval hierarchy, capacity constraints at a bottleneck institution, and a deliberate choice by some participants to settle late. Note what it does not contain: any suggestion that the parties lack the technical means to move faster. *[INFERENCE]*

The loan market does not settle slowly because it cannot. It settles slowly because settling requires the cooperation of parties who are not compensated for cooperating quickly, and because the market's own remedy for slowness exempts the largest single cause. *[INFERENCE]*

22.4 The reform that keeps being attempted

The trade association has campaigned on this repeatedly. It shifted delayed compensation to a requirements basis in 2016. *[FACT]* It introduced late-payment fees for loan trade transactions in May 2026. *[FACT]* Its 2026 programme focuses on standardisation, automation, and changes to credit-agreement provisions to shorten settlement, targeting consent tracking, agreement updates and lender communications — with no fixed settlement mandate. *[FACT]* Distressed settlement is singled out for change. *[FACT]*

The association's own diagnosis is telling: **"changing credit agreement provisions is necessary to shorten settlement times."** *[FACT]* That is a statement that the bottleneck lies in the contracts, not the systems. It is the most important sentence any standard-setter in this market has published, and Chapter 41 returns to it.

Delayed compensation and the problem of paying a moving syndicate

Because loans trade and settle slowly, the agent must pay principal and interest to a *moving* set of lenders. This produces a set of mechanics with no equivalent in most asset classes.

Delayed compensation, born in 1998, passes accruing interest and fees to the buyer net of the seller's cost of funds. *[FACT]*

Pre- and post-settlement cash-flow splits. Under the "payment on settlement date" convention the buyer pays the seller accrued and unpaid interest for the pre-T+7 period at closing; under the "settled without accrued interest" convention the buyer is credited for accrued interest for the window between T+7 and settlement. *[FACT]* These conventions literally allocate cash flows across the settlement boundary.

Lender-of-record tracking mid-trade. If a buyer was not a lender of record on the trade date, the seller must use commercially reasonable efforts to pass through all amendments and waivers arising between trade date and settlement; the agent pays interest on the payment date to the lender of record. *[FACT]*

The composite picture is this. **At any moment, a loan facility has a legal owner of record, an economic owner, and a set of in-flight trades in which the two differ, each with its own conventions for splitting cash and passing through votes.** *[INFERENCE]* The agent's register is authoritative on the first. Nobody is systematically authoritative on the second. Reconciliation between them is not an occasional exception; it is the permanent state of the market.

Distress, restructuring, and the par/distressed boundary

The distressed segment is a separate operating regime.


Documentation is heavier: distressed trades use a distinct trade confirmation and typically a negotiated purchase and sale agreement, with representations, "big-boy" letters, and diligence on the underlying claim — versus the lighter par confirmation. *[FACT]* Settlement is longer: T+20 against T+7. *[FACT]*

A subtlety worth noting: because participants increasingly traded distressed loans on par documents, the trade association folded distressed provisions into the par standard terms and conditions in its July 2023 revision. *[FACT]* That is a small but revealing episode — the standard-setter conforming the standard to what the market was doing anyway, rather than the market conforming to the standard.

When a credit deteriorates, a trade agreed on par documents may need to convert to distressed documents, or the reverse on recovery — a documentation-intensive reclassification changing representations, settlement window and delayed-compensation treatment. *[FACT]* No public figures exist on the frequency or dollar volume of such switches. *[UNKNOWN]*

A note on the length of this chapter. Distressed credit is deliberately out of scope. It has its own documentation regime, its own settlement window, its own reclassification machinery, its own claims-trading and bankruptcy-court overlay, and an operating problem that deserves to be analysed on its own terms rather than compressed into a chapter of a paper about the par market. A separate paper on distressed credit would be warranted if there is sufficient interest.

The duplication census



Here is the same fact, recomputed.

A borrower makes one interest payment. That single event is then independently processed by:

- the **agent**, which calculates the per-lender split, executes the payments, and records them;
- each **lender**, which records the receipt against its own position record and reconciles it against its own accrual;
- each lender's **fund administrator** or third-party administrator, which maintains a parallel set of books;
- each lender's **custodian**, where cash is held;
- the **CLO trustee**, for facilities held in CLOs, which reports the cash flow into the waterfall;

- the **fund accountant**, which strikes a net asset value incorporating it;
- the **pricing vendor**, which reflects the payment in its accrual and price data;
- and, on a delayed basis, any **rating agency** or **analytics provider** modelling the facility.

Each of these parties independently interprets the same credit agreement, ingests the same notice, performs the same or an equivalent calculation, and produces a record that must agree with the others. When two disagree, a break is opened and a human investigates.

The direct evidence for the burden is that the market has built multiple businesses to address it. An industry utility launched a reconciliation service in 2008 specifically because "lenders and banks reconciled loan positions independently of one another, a process prone to errors, miscalculations and delays," reaching roughly 2,300 funds. *[FACT]* A third-party administrator markets a "bank-loan data processing service eliminating manual agent-notice input" — evidence that manual agent-notice keying is a named, productized pain. *[VENDOR CLAIM]* A bank-founded platform launched a reconciliation module in January 2025 to electronically match lender and agent records and display the transaction notices causing the differences. *[VENDOR CLAIM]* A data-collaboration network operates a dedicated loan data management service on the basis that syndicated loans suffer from a "lack of transparency around agent data" where "absence of data reconciliation leads to incorrect payments and rework." *[VENDOR CLAIM]*

No published figure exists for the aggregate cost of this duplication.

[UNKNOWN] The nearest available anchors are a consultancy estimate that automation can cut loan-operations costs by roughly twenty to thirty per cent *[VENDOR CLAIM — COULD NOT BE VERIFIED FROM THE PRIMARY SOURCE, WHICH REPEATEDLY FAILED TO LOAD]* and the observation that syndicated loans represented approximately ten per cent of worldwide corporate-banking revenue in 2022 *[VENDOR CLAIM]*.


That the market cannot measure its own duplication cost is itself a finding,
and a curiously durable one. *[INFERENCE]*

FIGURE 3

One economic event, many institutional states

ILLUSTRATIVE EXAMPLE · SYNDICATED LOAN INTEREST PAYMENT

The cash moves once. The information does not. Each institution reconstructs the event from its own notices, systems, data and obligations — creating multiple, legitimate representations of the same economic reality.



Other touchpoints, not exhaustive. Legal and documentation · middle office · operations · tax · transfer agent · prime broker · data vendors · internal audit · board and committees.

DIVERGENCE, AND WHAT IT DOES AND DOES NOT MEAN

Different representations can diverge for legitimate timing, accounting, valuation or role-specific reasons — accrual basis, valuation date, allocation timing, accounting policy. Other differences are genuine breaks that require investigation. Structural duplication does not make every divergence legitimate, and the market has no shared way to tell the two apart quickly. [INFERENCE]

The duplication is structural, not accidental. Each institution has obligations that require its own representation of the event.

Institutions, systems and amounts are an illustrative composite. Estimated market cost of this duplication: [UNKNOWN]

The cash moves once; the information does not. Each institution reconstructs the same payment from its own notices, systems and obligations, and representations can legitimately diverge on timing, accounting basis and role — while other divergences are genuine breaks requiring investigation. No public source measures the aggregate cost of this duplication. *[UNKNOWN]*

Where the complexity actually lives

CATEGORY	CHARACTER & EXAMPLE	SUSCEPTIBILITY TO AI
1 • Calculation	Deterministic — interest accruals under compounded overnight rates with lookbacks. Given the terms and the facts, there is exactly one right answer.	None, and it should have none. A probabilistic system has no business computing a payment.
2 • Interpretation	Semantic, irreducibly ambiguous at the margins — reading a bespoke 200-page credit agreement to determine what the margin grid is. Two competent lawyers can disagree.	This is the category where the capability genuinely changed, and Part Nine examines it in detail.
3 • Coordination	Multi-party, sequential, dependent on other institutions' willingness to act — obtaining borrower consent to an assignment.	Marginal. It cannot make a counterparty decide faster, and it cannot supply a consent the counterparty declines to give.
4 • Data synchronisation	Mechanical, high-volume, structurally endless — reconciling positions between lender and agent.	Partial. Matching and exception classification are automatable; investigating a genuine discrepancy requires interpretation and often coordination.
5 • Institutional control	Political, contractual and economic, absolute in the sense that it does not yield to technical merit — who holds the register.	Zero.

This is the chapter on which the paper's later argument depends.

Everything described in Parts Two and Three can be sorted into five categories of complexity, and they have radically different properties.

Category 1: Calculation

Interest accruals under compounded overnight rates with lookbacks. Fee computations across a menu of fee types. Amortization schedules. PIK capitalisation and its compounding effects. Waterfall priority-of-payments.

Delayed-compensation true-ups. Pre- and post-settlement cash splits.

Character. Deterministic. Given the terms and the facts, there is exactly one right answer.

Difficulty. Moderate and entirely tractable. This category has been solvable by software since the 1990s and is solved by the incumbent platforms.

Susceptibility to AI. None, and it should have none. A probabilistic system has no business computing a payment. *[INFERENCE]* Chapter 65 develops this.

Category 2: Interpretation

Reading a bespoke 200-page credit agreement to determine what the margin grid is, when the prepayment penalty applies, what triggers the excess-cash-flow sweep, whether this event is a permitted acquisition, what the covenant actually requires. Reading an unstructured notice to determine what event it describes and what the sender means by "facility fee."

Character. Semantic and irreducibly ambiguous at the margins. Two competent lawyers can disagree.

Difficulty. High, and it has been the market's binding constraint. The evidence that this is the binding constraint is that the commercially important terms "remain as bespoke to the specific transaction as ever" and the model forms "are rarely used as a starting draft." *[FACT]* A rules engine requires someone to pre-encode every covenant, pricing grid, basket and event trigger as structured configuration; when each agreement is custom-negotiated, that encoding is manual, per-deal and brittle — which is the documented reason operations shops fall back to spreadsheets. *[FACT]*

— THE FALLBACK TO SPREADSHEETS IS DOCUMENTED; THE CAUSAL EXPLANATION IS INFERENCE] [INFERENCE]

Susceptibility to AI. This is the category where the capability genuinely changed, and Part Nine examines it in detail.

Category 3: Coordination

Obtaining borrower consent to an assignment. Obtaining agent acknowledgment. Collecting a hundred lender votes to a threshold by a deadline. Getting a signature from an authorised officer. Scheduling a bank meeting. Chasing a counterparty who has not returned a document.

Character. Multi-party, sequential, dependent on other institutions' willingness and capacity to act.

Difficulty. High, and — this is the essential point — **not a technical problem.** The reason a consent takes three days is that a person at another institution has to decide to look at it. Every one of the eight settlement-friction items enumerated in Chapter 22.3 falls in this category, which is why that list is worth re-reading here rather than restating: it is not a list of technology gaps, it is a list of decisions other people have to make. *[INFERENCE]*

Susceptibility to AI. Marginal. Software can reduce the latency of *transmitting* a request and *tracking* its status. It cannot make a counterparty decide faster, and it cannot supply a consent the counterparty declines to give. *[INFERENCE]* Any thesis whose economics depend on compressing this category should be examined with hostility.

Category 4: Data synchronisation

Reconciling positions between lender and agent. Reconciling cash. Detecting and investigating breaks. Maintaining consistency across the agent, the lender, the administrator, the custodian and the trustee. Resolving identifier mismatches across four competing identifier schemes.

Character. Mechanical, high-volume, and structurally endless because it is generated by the market's architecture rather than by anyone's error.

Difficulty. Moderate technically, high organisationally.

Susceptibility to AI. Partial. Matching and exception classification are automatable; investigation of a genuine discrepancy requires interpretation (Category 2) and often coordination (Category 3). *[INFERENCE]*

This is the category the existing utilities and networks address, and it is the most contested competitive space in the market. *[INFERENCE]*

Category 5: Institutional control

Who holds the register. Who decides whether an assignment is processed today or Thursday. Who sets and waives the assignment fee. Who decides whether the bank buys software or builds it. Who decides whether a consortium is formed. Who bears liability if a payment goes wrong.

Character. Political, contractual and economic.

Difficulty. Absolute, in the sense that it does not yield to any amount of technical merit.

Susceptibility to AI. Zero. *[INFERENCE]*

26.1 The argument

With the decomposition in hand, the market's history becomes legible.

Across the initiatives examined in this paper, the technologies that succeeded addressed Category 1 or Category 4. The incumbent servicing platforms solved calculation and record-keeping. The settlement platforms and reconciliation utilities addressed synchronisation. *[INFERENCE]*

In every material failure identified here, the failure was in Category 3 or Category 5. The messaging standard failed because adoption required coordinated action by parties who each bore cost for collective benefit. The settlement-time reforms failed because the friction is consent, not compute. The distributed-ledger ventures failed, on the documented evidence, because of "legal enablement, network coordination, and DLT being outcompeted by ordinary centralized software" — coordination and control, not technology. *[FACT] [INFERENCE]* Part Six sets out the full survey on which this generalisation rests, and the reader should treat it as a pattern across six waves rather than a law.

Category 2 was, until recently, simply written off. No one seriously proposed to read every credit agreement, because reading every credit agreement was not economically possible. The market's response was to have humans read the ones that mattered and to run the rest on spreadsheets. *[INFERENCE]*

That is the setup for the central question. If something has changed, the change is in Category 2 — and the question the rest of this paper must answer is whether a change in Category 2 is sufficient to move a market whose failures have historically been in Categories 3 and 5. *[INFERENCE]*

The reader should hold that formulation. It is the fairest and most demanding statement of the problem this paper can offer. It is the standard against which Part Eleven tests the strategic hypothesis, and Part Twelve asks the same question of the industry itself.

IV

PART FOUR

Incumbency

Why Loan IQ and ACBS won, and stayed won

There is a version of this market's story in which the incumbent servicing platforms are bad software that survives through inertia, and the opportunity is simply to build better software. That version is wrong in a way that matters, and this part is an attempt to say precisely how.

The incumbents did not win by accident and are not held in place by inertia alone. They won a genuine technical race in the 1990s, at a moment when winning it was hard. They then accumulated a moat that has almost nothing to do with product quality and cannot be attacked with product quality. Understanding the composition of that moat — layer by layer, with each layer's strength assessed independently — is the difference between a strategy that has a chance and one that does not.

This part also observes an evidentiary discipline that the surrounding literature generally does not. The public record on the incumbent platforms is thin, and much of what circulates about them is either vendor marketing or practitioner folklore. Where the record is thin, this part says so.

The core-system race

By the early 1990s the market had a problem that its participants could feel. Volumes were growing, syndicates were getting larger, secondary trading was becoming real, and the servicing function — the register, the accruals, the notices, the payments — was running on a combination of mainframe programs, spreadsheets and paper that did not scale.

Two systems emerged to solve it.

ACBS — the Advanced Commercial Banking System — had a design team form in 1991, a prototype within a year, and a first customer in 1993. It was built on a fourth-generation application platform for the midrange environment. *[VENDOR CLAIM — THE DATES AND PLATFORM COME FROM VENDOR AND PARTNER MATERIALS]* It was marketed as a high-end origination, syndication, servicing, trading and settlement system for corporate and large syndicated loans. Its vendor at one point claimed that "over 30 per cent of syndicated loans worldwide are processed by ACBS." *[VENDOR CLAIM]*

Loan IQ originated as the flagship product of **IQ Financial Systems**, a New York commercial-lending and capital-markets software house that was a majority-held subsidiary of **Deutsche Bank**. *[FACT]* It was acquired by Misys for **\$40 million in cash**, announced around January 2004; at acquisition IQ Financial Systems had approximately \$22 million of revenue for its 2003 financial year and served roughly twenty-nine clients, predominantly in North America. *[FACT]*

A correction is required here because the error is widespread. **Loan IQ did not originate from "Midas."** Midas is a separate core-banking product line first developed in the 1970s and is not the ancestor of Loan IQ. *[FACT]* Any document tracing Loan IQ to Midas is repeating a mistake.

The subsequent ownership chain: Vista Equity Partners took Misys private in March 2012 for approximately £1.3 billion; Vista acquired the Canadian financial-technology firm D+H in 2017 at C\$25.50 per share including assumed debt for a total enterprise value of approximately C\$4.8 billion, announced 13 March and completed 13 June 2017; and the two were merged in June 2017 to form Finastra, described at formation as the third-largest financial-technology company in the world. *[FACT]* The frequently cited "\$3.6 billion" figure for D+H is the US-dollar equivalent of the C\$4.8 billion take-private, not a separate number. *[FACT]*

ACBS's chain is equally often garbled and is worth stating correctly. **ACBS did not reach FIS through SunGard.** It came via the acquisition of ALLTEL Information Services' financial-services division by Fidelity National Financial, announced 29 January 2003 and completed 1 April 2003 for approximately \$1.05 billion; the acquired business was renamed Fidelity Information Services. *[FACT]* FIS's separate acquisition of SunGard, announced 12 August 2015 at approximately \$9.1 billion including assumed debt, is a real event that carried SunGard's capital-markets and asset-management assets into FIS — but it did not carry ACBS, which was already there. *[FACT]* There is no evidence SunGard ever owned ACBS. *[FACT]* The product is today branded FIS Commercial Loan Servicing. *[FACT]*

Two further name collisions in this lineage are worth clearing, because they recur in secondary literature. The "AFS" that SunGard acquired in December 2006 was **Automated Financial Systems Corporation**, a brokerage and wealth front-office vendor absorbed into a wealth-management product — not the Exton, Pennsylvania commercial-lending firm **Automated Financial Systems, Inc.**, developer of AFSVision, which remained independent until its acquisition by a private-equity firm announced in February 2026. *[FACT]* And "ACBS" stands for *Advanced Commercial Banking System*; "Automated" is a common mislabel. *[FACT]*

27.1 What they actually won

The systems that won this race did something genuinely difficult. They modelled a facility structure capable of representing multi-currency, multi-tranche, multi-branch lending with a syndicate whose composition changes over time; they executed the full menu of credit events against those structures; they maintained a register that could survive assignment, participation and restructuring; and they generated the notices and payments that the market's operational reality requires.

That is not a trivial data model. It is arguably the hardest data model in commercial banking, and Chapter 30 argues that the difficulty of the model is itself one layer of the moat. *[INFERENCE]*

The market's judgment on the outcome is expressed in migration flows. The documented flow is **one-directional: ACBS to Loan IQ**. *[FACT]* Reporting around one major bank's migration from ACBS to Loan IQ noted that most European ACBS customers had migrated over several years, leaving "just a handful" on the continent — while the incumbent vendor countered that it had "lost only a couple of clients over the past ten years." *[FACT — REPORTED;*

THE TWO CLAIMS ARE IN TENSION AND THE REPORTING IS FROM A SINGLE TRADE OUTLET WHOSE PAGE BLOCKS AUTOMATED RETRIEVAL, SO THE SPECIFIC "HANDFUL" CHARACTERISATION SHOULD BE TREATED AS REPORTED RATHER THAN INDEPENDENTLY CONFIRMED]

There is no public case of any bank migrating off Loan IQ to a competitor.

[FACT — AN ABSENCE, ESTABLISHED BY SEARCH RATHER THAN BY A POSITIVE SOURCE, AND THEREFORE AN ABSENCE OF EVIDENCE RATHER THAN EVIDENCE OF ABSENCE; BUT A CONSPICUOUS ONE]

Loan IQ: what is known, what is claimed, what is unknown

Discipline is required here because almost everything widely believed about this platform comes from its vendor.

28.1 What the vendor claims

The following are **Finastra's own marketing claims**, published on its website and in product brochures, and are **not independently verified**:

- Loan IQ is used by **21 of the top 25** syndicated lenders. *[VENDOR CLAIM]*
- Loan IQ processes approximately **70 per cent of the world's syndicated loan volume**. (A partner brochure has used 75 per cent.) *[VENDOR CLAIM]*
- **9 of the top 10** global agent banks use Loan IQ. *[VENDOR CLAIM]*
- Approximately **\$7.5 trillion** of global loan facilities under management. *[VENDOR CLAIM]*
- Loan IQ is "the most trusted commercial and syndicated loan servicing platform" and "the market's preeminent loan servicing platform." *[VENDOR CLAIM]*

No independent analyst dataset quantifying Loan IQ's syndicated-loan share could be located in public sources. *[UNKNOWN]* This is a significant gap. The "70 per cent" figure is the single most-repeated number in every strategic document about this market, it is repeated in press coverage and in third-party analyses, and its only source is the vendor. A paper that cites it as fact is citing marketing.

That said, the *directional* claim is corroborated from an unexpected quarter. A data-collaboration network's integration announcement describes its partner's lender-data product as powering "more than 70 per cent of the world's syndicated loans" — but that is the same vendor's figure quoted by a partner, not an independent measurement. *[VENDOR CLAIM]* The genuinely independent corroboration is weaker and more circumstantial: the one-

directional migration flow, the absence of any documented off-ramp, and the fact that market participants' own modernization efforts are designed around Loan IQ as the assumed source system. *[INFERENCE]*

Related firm-wide figures: approximately 9,000-plus customers across roughly 130 countries, 48 of the top 50 banks, and roughly \$1.9 billion of revenue. *[VENDOR CLAIM]* A widely repeated "8,000+ customers" figure does not match current vendor materials and appears dated. *[FACT]*

28.2 What is verified

- The product's origin, price, and acquisition chain (Chapter 27). *[FACT]*
- Modernization is delivered through **Loan IQ Nexus**, an integration and interoperability layer with modules for build, transact, maintain and connect, offering REST interfaces, cloud-enabled deployment and, per vendor materials, AI-assisted onboarding — with the vendor managing translation and integration to the core. *[FACT THAT THE PRODUCTS EXIST; VENDOR CLAIM AS TO THEIR CAPABILITIES]*
- The platform is available as software-as-a-service or on premises. *[FACT]*
- Pricing is **not public**; it is sold on a contact-vendor basis with no published range, no free trial and no free version. *[FACT]*
- The vendor published a market commentary titled "**Busting the Loan IQ myth**" in 2021. The myth it rebuts is that Loan IQ is only for complex syndicated loans and is too heavy for high-volume small-business and bilateral lending. *[FACT]*
- A named migration flow exists in the market for ACBS-to-Loan-IQ data migration specialists, evidenced by active job postings at multiple staffing and consulting firms. *[FACT]*

That "myth-busting" document deserves a moment. A vendor publishing a defence of its flagship product against a specific market perception is *prima facie* evidence that the perception exists and that the vendor considers it commercially material. *[INFERENCE]* It is not evidence that the perception is correct. It is, however, the clearest public signal available that the platform is understood in its market as heavyweight and narrow.

28.3 What is inferred

- The architecture has substantial legacy heritage. One research pass inferred an original stack of a client-server development environment over a relational database of that era, from job postings for migration developers on Loan IQ engagements and from an architecture document referencing a database product of the period. *[INFERENCE — THIS IS A CHAIN OF CIRCUMSTANTIAL INFERENCE FROM RECRUITMENT MATERIAL, NOT A TECHNICAL FACT, AND IT SHOULD NOT BE ASSERTED AS ONE]*
- Modernization is **incremental** — cloud enablement and an interface layer around a legacy core — rather than a rebuild. *[INFERENCE, FROM THE STRUCTURE OF THE NEXUS PRODUCT LINE AS DESCRIBED IN VENDOR MATERIALS]*
- Migration off is expensive and long. Reporting frames loan-servicing platform migrations generally as "a lengthy and expensive undertaking." *[FACT]* The inference is that this applies to Loan IQ specifically. *[INFERENCE]*

28.4 What is unknown

The following could not be established from public sources and should be treated as genuinely open:

- Current release cadence, version, and technical stack in any specific detail. *[UNKNOWN]*
- Pricing model, contract structure, and the economics of a typical customer relationship. *[UNKNOWN]*
- Active customer count beyond the vendor's own claims. *[UNKNOWN]*
- Any independent market-share measurement. *[UNKNOWN]*
- Whether the vendor's DLT product (Chapter 45) is still operating, was folded into the interface layer, or was retired. *[UNKNOWN]*
- Whether any specific pricing change was made that customers resented. This claim circulates widely in practitioner discussion; **no external reporting of any such change could be found**. Searches across trade press, review sites and pricing aggregators returned only generic descriptions of licence-plus-maintenance and subscription models, with no documented grievance, fee-hike event or re-pricing story. *[UNKNOWN]* It should not be asserted.

28.5 The satisfaction evidence

The public review corpus is thin and middling rather than damning.

A leading enterprise-software review platform shows the vendor at 3.3 stars on only three reviews in commercial loan origination. Representative criticisms are **slow support resolution** — "problems requiring support, with resolutions not as quick as needed, often bleeding over reporting periods" — and **limited customization** — users "cannot customize the product fully... wishing for more customization options." Positives cite single-platform consolidation and reduced error rates. *[FACT — AS TO WHAT THE REVIEWS SAY; N=3 CANNOT CARRY ANY THESIS]* A second review platform rates the vendor 8.0 out of 10 and ranks it first in loan management, while repeating the same support complaint nearly verbatim. *[FACT]*

Enterprise loan-servicing buyers do not leave public reviews at volume, and three reviews is not data. *[INFERENCE]* The recurring, independently echoed complaint is support responsiveness and customization rigidity — irritations, not catastrophes. Anyone building a case on customer hatred of this platform is building on a foundation the public record does not supply. *[INFERENCE]*

ACBS and the other side of the field

FIS's position is worth stating because it is the market's only structural counterexample and because it is instructive about what happens to the loser of a core-system race.

FIS is a large public company — 2024 revenue of approximately \$10.1 billion, roughly 45,800 employees, with banking solutions and capital-market solutions segments. *[FACT]* Its commercial loan servicing product carries the ACBS heritage described in Chapter 27, alongside a syndication product, and FIS also operates a middle-office services business positioned as lowering "the barrier to the credit market." *[FACT/VENDOR CLAIM]*

Three observations.

A losing core system does not disappear; it becomes a migration source.

The market's specialist consultancies staff ACBS-to-Loan-IQ migration projects, and one managed-services firm explicitly offers ACBS-to-Loan-IQ migration and reconciliation. *[FACT]* The loser's installed base becomes the winner's pipeline, slowly, over decades.

The vendor disputes the narrative, and the dispute cannot be resolved from public sources. FIS's counter — that it lost "only a couple of clients over the past ten years," one of them via an acquisition rather than a competitive loss — is a claim, and so is the trade-press characterisation that "just a handful" of European ACBS customers remain. *[FACT THAT BOTH CLAIMS WERE MADE; UNKNOWN AS TO WHICH IS CLOSER TO CORRECT]* Neither party publishes a client count, the reporting sits behind a page that refuses automated retrieval, and no independent census of loan-servicing installations exists. What can be said with confidence is narrower and still useful: **the documented migrations run in one direction, and a market for ACBS-to-Loan-IQ migration specialists exists and is actively recruiting, which is not something that happens around a stationary installed base.**

[FACT] [INFERENCE]

There is a second reading of the same evidence that deserves airing, because it cuts against this paper's general framing. If FIS is right that it has lost only a couple of clients in a decade, then the migration flow is a trickle rather than an exodus, and the correct conclusion is not that ACBS lost the race but that **both cores are extraordinarily sticky and the market simply does not move.** *[HYPOTHESIS]* That reading is consistent with everything in Chapters 30 to 32 and would, if true, strengthen rather than weaken the paper's argument about switching costs.

FIS is a competitor with its own services arm. Its middle-office business puts it in the same category as the platform-owning administrators identified in Chapter 11: a firm whose services revenue and platform revenue are complementary and whose interest in a third-party coexistence layer is therefore negative. *[INFERENCE]*

Beyond the two incumbents, the field of servicing platforms includes an independent real-time commercial lending platform acquired by a private-equity firm in February 2026; cloud-native origination and lifecycle platforms serving relationship and mid-market lending; a major core-banking vendor's corporate lending module; and several payments and core-banking companies whose complex-commercial-lending route is via partnership rather than an owned syndications engine. *[FACT]* The comparison material indicates these largely serve different segments — commercial and consumer origination, embedded finance, mid-market — rather than displacing the incumbents in complex top-tier syndicated servicing. *[INFERENCE]*

The anatomy of the moat

LAYER	WHAT IT IS	STRENGTH
1 • Data model	Multi-currency, multi-tranche facility structure. The layer challengers attack, and the one that matters least.	Moderate
2 • Books and records	The system of record feeding the ledger, risk and regulatory reporting.	High
3 • Integration surface	Payments, treasury, GL, risk, KYC, reporting — each connection built once, maintained by someone.	High, grows with tenure
4 • Operational muscle memory	Staff who know the screens and workarounds; one-to-two-year training against twenty-to-seventy-per-cent attrition.	Very high, increasing
5 • Career risk	Nobody was promoted for a migration that went well.	Absolute

The incumbent position is held by five distinct layers. They are usually discussed as one thing. They are not one thing, they have very different strengths, and a strategy that attacks the wrong layer will fail regardless of its merits.

Layer 1 — The data model

The facility structure described in Chapter 27.1 is genuinely hard to build and harder to build correctly. Vendor material describes a data model natively supporting multi-currency facilities, complex fee structures, syndicate position management and secondary-market trading. *[VENDOR CLAIM]* The claim is self-serving and probably true. *[INFERENCE]*

Strength: moderate. This is the layer most vulnerable to a competent modern rebuild, and — across the challengers surveyed in Part Six — it is the layer challengers have generally attacked. It is also the layer that matters least, because a bank does not stay on a platform for three decades because of its data model. *[INFERENCE]*

Layer 2 — Regulatory books and records

The servicing platform is the system of record for the bank's loan book. It feeds the general ledger, risk systems, regulatory reporting and audit. Replacing it means re-establishing that entire chain of evidence to the satisfaction of internal audit, external audit and the supervisor.

Strength: high. This is not addressed by any product feature. *[INFERENCE]*

Layer 3 — Integration surface

The platform is connected to payments, treasury, general ledger, credit risk, customer data, document management, KYC and reporting. Each connection was built once and is maintained by someone.

Strength: high, and it grows with tenure. *[INFERENCE]*

Layer 4 — Operational muscle memory

Hundreds of operations staff know the platform's screens, workarounds and quirks. Procedures, controls, training materials and audit evidence are written against it. Chapter 66's evidence on training time — one to two years to productivity in loan operations *[VENDOR CLAIM]* — is the measure of this layer: replacing the platform means re-training a workforce whose training already takes years, while the workforce is turning over at twenty to seventy per cent. *[VENDOR CLAIM] [INFERENCE]*

Strength: very high, and increasing as the workforce hollows out. This is the counterintuitive result: the knowledge cliff makes migration harder, not easier, because migration is precisely the activity that requires deep platform knowledge. *[INFERENCE]*

Layer 5 — Career risk

Nobody has ever been promoted for a loan-servicing migration that went well. Several have presumably suffered for ones that went badly. The decision to replace a core servicing platform is taken by an executive who

bears asymmetric personal risk, and the literature's framing of such migrations as "lengthy and expensive" is the language in which that risk is expressed. [FACT] [INFERENCE]

Strength: absolute, in the sense that it is not addressed by evidence. [INFERENCE]

30.1 The implication

Layers 2 through 5 are untouched by product quality. A challenger with a strictly better system faces the same integration re-work, the same audit re-establishment, the same retraining, and the same career risk as a challenger with a worse one. [INFERENCE]

There is a sixth element that is not a layer of the bank at all, and leaving it out understates the position. **A liquid external market exists for people who can implement, migrate to, test and support the incumbent platforms, and no equivalent market exists for a challenger** (Chapter 13).

[INFERENCE] That asymmetry does not create switching cost by itself; it multiplies Layers 3 and 4, because the bank that changes core has to re-acquire integration and operating knowledge in a market that has not yet formed for the alternative — at the moment it needs depth most. The moat is therefore better described as an ecosystem than as a product: software, configuration, internal knowledge, external expertise, procedures, contractors, testing infrastructure and career capital, accreted over three decades. [INFERENCE]

This is the single most important strategic fact in the paper about the incumbent position, and it has a direct corollary: any strategy that requires a bank to replace its core servicing platform is competing on four dimensions where product quality is largely beside the point. [INFERENCE] A better system reduces the *ongoing* cost of Layer 4 once the migration is complete; it does nothing to reduce the cost of getting there, and the cost of getting there is what the decision turns on.

The distinction matters because it changes the competitive question. A better servicing product is competing not only with another product. It is competing with a system of record, integrations, controls, procedures, trained people, audit evidence and regulatory confidence accumulated around that product over decades. [INFERENCE]

The incumbent's deepest advantage is therefore institutional rather than technological. Product quality can narrow a capability gap. It cannot, by itself, reproduce the institutional state surrounding the incumbent.

It also has a less obvious corollary that cuts the other way, and Part Eleven will lean on it. If the moat is composed of switching costs rather than product love, then the moat protects against *replacement* and does nothing against *addition*. A product that does not ask the bank to replace anything faces none of Layers 2 through 5. [INFERENCE] Whether that is a genuine opening or a consolation prize is the question of Part Eleven.

FIGURE 4

The installed core is institution-specific

The same product has evolved into different institutional infrastructure at different banks — and often on different release paths.

SEVEN DIMENSIONS OF DIVERGENCE

1 • Common product One platform from the vendor	2 • Different versions Banks are on different releases
3 • Different configurations Different modules, different setups	4 • Different customizations Institution-specific code and workflows
5 • Different integrations Unique internal and external connections	6 • Different controls Bank-specific governance and reporting
7 • Different operating models Different people, processes and risk appetites	


SAME PRODUCT, DIFFERENT BANKS

	BANK A • GLOBAL MONEY CENTRE	BANK B • REGIONAL, SPECIALIZED	BANK C • INTERNATIONAL
Version / release	A more recent release, selectively upgraded	An earlier release, upgrades deferred	An older release, heavily customized
Modules in use	A broad module footprint across servicing, accounting, trading and risk	A narrower footprint; some functions held in other systems	A different subset again, with other functions elsewhere
Customizations	Extensive — many bank-specific extensions	Moderate — some bank-specific extensions	Extensive — bank-specific workflows
Key integrations	Core banking, SWIFT, general ledger, data warehouse, regulatory reporting, client portals	Core banking, ACH / Fed, external credit risk, data warehouse, investor reporting	Core banking, local payment systems, local regulatory, custom finance / GL, client reporting
Upgrade experience	High complexity — months of testing, interface impacts	Challenging — often deferred	Very high complexity — major programme, many dependencies

Release paths diverge accordingly: institutions sit on different releases, each with its own history of major upgrades, minor updates and deferrals. No release numbers or dates are asserted here.

FIGURE 4, CONTINUED

INSIDE ONE BANK — A DECADE OF ACCUMULATION AROUND THE CORE



Key observations. No two implementations are alike in modules, configuration, extensions or integrations · banks sit on different releases, and upgrades are known to be complex, lengthy and risky · SDK-based extensions and bank-specific logic are common at major lenders · the core sits inside a much larger network of systems, controls, people and external participants · replacement or major upgrade can disrupt critical operations, reporting and regulatory processes.

Years of interfaces, customizations, controls, people and processes: not just software, an institution. A replacement therefore competes with an institution's accumulated architecture, not with the vendor's product.

Bank profiles are illustrative composites of market patterns, not a survey of named institutions. No versions, dates or extension counts are asserted.

The same vendor platform exists at different banks as different institutional infrastructure: different releases, configurations, extensions, integrations, controls and operating models. This is the layered moat of Chapter 30 — a challenger competes with what the institution accumulated around the software, not with the software alone. Bank profiles shown are illustrative composites. *[INFERENCE]*

The switching-cost arithmetic

The public evidence on migration cost is qualitative rather than quantitative, and it is worth assembling what exists.

The characterisation. "Moving to a new lending system is a lengthy and expensive undertaking." *[FACT]* Structural drivers named across sources: multi-year data migration, integration with general ledger, risk and regulatory systems, operational-disruption risk, and loss of accustomed functionality. *[FACT]*

The scale of what is being migrated. Some banks service more than 300,000 loans on the incumbent platform. *[VENDOR CLAIM]*

The one-directional flow. ACBS to Loan IQ is documented; the reverse is not. *[FACT]*

The absence. No public case exists of a bank migrating off Loan IQ. *[FACT — AS AN ABSENCE]*

No published cost figure for a loan-servicing migration could be located. *[UNKNOWN]* This is a real gap. A strategic argument that turns on migration economics is, on the public record, turning on an unmeasured quantity.

What can be said is structural. A migration of this kind involves: extracting and transforming the register and all historical positions for hundreds of thousands of facilities; re-establishing every integration in Layer 3; running the old and new systems in parallel long enough for audit to accept the new one; retraining an operations population; rewriting procedures and controls; and re-obtaining whatever supervisory comfort the institution's regulator expects. Each of those is a multi-quarter workstream and several are prerequisites for others. *[INFERENCE]*

The honest conclusion is that the switching cost is large enough that no bank has publicly borne it, and that the market therefore has no observed price for the thing everyone assumes is expensive. *[INFERENCE]*

The dissatisfaction–switching gap

The best-evidenced part of the incumbency story is not that customers love the platform. It is that whether they love it is nearly irrelevant to whether they leave.

The nearest public analogue is the American Bankers Association's Core Platforms Survey, conducted in late 2024 and released in February 2025. It found only fifty-three per cent of banks satisfied with their core vendor and roughly thirty-five per cent dissatisfied — yet **only about one in five, nineteen per cent, expected to switch at renewal**, with sixty-nine per cent likely to stay. *[FACT]* All four figures belong to that single survey; the 2024/2025 dating that circulates reflects the gap between fieldwork and publication, not two different studies.

The survey's most useful finding for present purposes is a curve rather than a level: **"satisfaction declines the longer a bank operates on their current core platform... dropping to the lowest level as the bank approaches contract renewal."** *[FACT]*

Three things follow.

Dissatisfaction is real and does not convert. The gap between thirty-five per cent dissatisfied and nineteen per cent expecting to switch is the switching cost, expressed as a survey statistic. *[INFERENCE]*

The dissatisfaction is back-loaded. It peaks at the moment of maximum commercial leverage, which is also the only moment at which a change is contractually contemplated. *[FACT]*

The analogue is imperfect and must be labelled as such. This is a survey of *core banking platforms*, not loan-servicing platforms, and the customer populations differ — core platform buyers include many small institutions, whereas Loan IQ's base is concentrated among large agent banks. *[FACT]* Whether the same curve holds for syndicated loan servicing is **unknown**,

and the paper does not assert it. *[UNKNOWN]* What the analogue establishes is that the pattern — high dissatisfaction, low switching, satisfaction worst at renewal — is a documented feature of enterprise banking software generally, not a story invented about one vendor. *[INFERENCE]*

The private-equity overlay

The ownership structure of the market's software is unusual and consequential, and the analysis of it is frequently overstated.

33.1 The facts

One sponsor owns both sides. Vista Equity Partners owns Finastra, and therefore Loan IQ — the sell-side syndicated servicing backbone. *[FACT]* Vista also controls Allvue Systems, formed in September 2019 when Vista acquired AltaReturn and merged it with Black Mountain Systems — infrastructure for the private-credit and CLO managers who increasingly compete with and buy from banks. *[FACT]* Allvue's own registration statement describes it as "controlled by certain affiliates of Vista Equity Partners." *[FACT]* Allvue's footprint is material: reporting on Vista's sale exploration states Allvue tracks more than \$7 trillion of assets under management on its platform and serves fifty of the top hundred CLO managers, with more than \$200 million of annual recurring revenue growing at more than fifteen per cent. *[VENDOR CLAIM / PRESS]*

The leverage is extreme. Finastra is among the most heavily levered software credits in the market. S&P has put adjusted leverage at approximately **16 times**, or roughly 8.5 times excluding preferred equity that the rating agency treats as debt-like. *[FACT]* The chronology: a proposed roughly \$6 billion private-credit refinancing met lender resistance in 2023, with net leverage reported around 9.97 times as of May 2023; in September 2023 Finastra closed **\$5.3 billion** — a \$4.8 billion unitranche term loan plus a \$500 million revolver at SOFR plus 7.25 per cent at 98 original issue discount — described as the largest US private credit transaction to date, with Vista contributing roughly \$1 billion of preferred equity. *[FACT]* In **April 2025** an attempt to refinance more than \$5 billion into the broadly syndicated market **collapsed after weeks of negotiations**. *[FACT]* It was

revived and completed by July 2025 at roughly \$3.6 billion equivalent, reported oversubscribed "despite leverage concerns" and shaving roughly 300 basis points. *[FACT]*

The portfolio is being pruned. Finastra agreed to sell its Treasury and Capital Markets division — announced 19 May 2025, covering products serving 340-plus institutions — to a private-equity buyer. *[FACT]* Under its chief executive, four divestitures have been completed since January 2025, including a mortgage-technology business, a universal-banking business, and a student-loan servicing business. *[FACT]* Vista is separately exploring a sale of Allvue at up to approximately **\$3 billion** including debt, with advisers engaged, as of August 2026. *[FACT]*

Security incidents. In November 2024 Finastra detected unusual activity on an internally hosted secure file transfer platform; an attacker used compromised credentials with access at intervals between 31 October and 8 November 2024, and a threat actor advertised approximately 400GB of allegedly stolen data. *[FACT]* **This was not a ransomware attack** — no malware was deployed, no lateral movement was found, and the company stated there was no direct impact on customer operations. *[FACT]* The 400GB figure and the specific exfiltration method are the actor's and secondary reporting's claims, not company-confirmed. *[FACT]* A further leak-site campaign in August 2026 named Finastra with a reported 400GB-plus of stolen data including client data; **the exact 2026 scope and whether loan-servicing data is implicated is not firmly established across independent outlets and should be treated as provisional.** *[UNKNOWN]* The systemic-risk framing is real: the vendor serves a large share of the world's major banks and the compromised system handled files containing wire and transfer instructions. *[FACT]*

33.2 What the overlay does and does not establish

It establishes concentration and financial constraint. One sponsor owning the dominant sell-side servicing platform and a leading buy-side platform is a genuine market-structure fact, and its implications appear to be

undiscussed in any public analysis — this paper found none. *[UNKNOWN]* The reflexive irony is documented rather than editorial: Finastra's own 2023 debt was refinanced by the private-credit industry that Allvue serves. *[FACT]*

It does not establish deliberate under-investment. No independent analyst report brands Finastra an under-investing shop coasting on lock-in.

[UNKNOWN] The available frame is circumstantial and mixed. The general private-equity software playbook — acquire sticky mission-critical software, raise prices five to eight per cent annually where switching costs exceed the increase, shift customers to multi-year contracts, cut costs — is well described in industry commentary and is real. *[VENDOR CLAIM / INDUSTRY COMMENTARY]* But its specific application to Loan IQ pricing is **not externally corroborated** (Chapter 28.4), and there is contrary evidence: the vendor did fund a cloud and open-platform programme with an application marketplace on a major public cloud, reporting hundreds of registered participants. *[VENDOR CLAIM — THE VENDOR'S OWN PRESS]*


The academic evidence is also weaker than the rhetoric. The most-cited study of private-equity effects on products found buyout targets grew sales roughly fifty per cent mainly via new products and geographic expansion, with only about one per cent price increases on existing products — which does not support a simple price-gouging story, though the study concerns consumer products rather than enterprise software. *[FACT]* Broader work finds employment falls around thirteen per cent over two years after take-privates. *[FACT]* None of it is specific to banking software. *[FACT]*

Most importantly: the overlay does not predict divestiture of the asset in question. This is the single most consequential correction to the conventional strategic narrative about this market, and it is recent.

Finastra is explicitly keeping Loan IQ and pointing it at private credit, with a dedicated private-credit positioning of the platform and the Nexus interface layer. *[FACT]* Lending and payments bookings were reported up twenty per cent for 2026. *[VENDOR CLAIM]* The company is pruning *around* Loan IQ, not selling it.

Any strategy premised on the incumbent platform becoming available, distracted, or abandoned should be revised. *[INFERENCE]* The incumbent is modernising and re-aiming, not vacating.

The talent diaspora



A quieter fact, with real strategic content.

The originating firm, IQ Financial Systems, was absorbed into Misys for \$40 million in 2004 — a small price for what became the market's dominant platform. *[FACT]* The originating architects are not publicly named. *[UNKNOWN]*

The platform's implementation expertise has since dispersed. A New York corporate- and commercial-lending consultancy founded in 2013 and specialising in Loan IQ implementation, alongside loan servicing, settlement, liquidity and automation work, was acquired by a global consultancy in October 2018 — folding a Loan IQ implementation brain-trust into a systems integrator. *[FACT]* A specialist boutique now operates as a Loan IQ advisory, implementation and support firm staffed by platform alumni, including a managing partner who previously led the platform's sales and global services delivery at that global consultancy and was a managing partner at the acquired firm. *[FACT]* The boutique runs a "time share" model leveraging platform-knowledgeable independent contractors.

[VENDOR CLAIM]

A standalone consultancy built entirely from one platform's know-how implies that the platform's expertise increasingly lives outside its vendor. *[INFERENCE]* That is the knowledge cliff of Chapter 66 reaching the vendor itself.

Adjacent to this, the loan-technology startup field has drawn leadership from the market's data and settlement infrastructure rather than from Finastra. **LendOS**, the most prominent private-credit lifecycle entrant, was founded in 2022 by **John Olesky**, formerly global co-head of loan platforms at IHS Markit — a franchise covering ClearPar settlement, deal management and notice management — with **Josh Herrera** as president. It launched a software platform in 2025 and announced a Series A led by Blackstone Innovations Investments in September 2025, the round's size undisclosed. *[FACT]* That pedigree is the **loan-trading and settlement stack**, not the servicing core — adjacent expertise rather than incumbent expertise. *[FACT]*

The strategic reading, which Part Thirteen develops: the people who know how to implement and operate the incumbent platform are now distributed across consultancies and boutiques whose economics depend on that platform continuing to exist, and who already resell third-party technology into their installed base. *[FACT]* They are a channel for anything additive and an enemy of anything replacing. *[INFERENCE]*

What the incumbency case does not prove

This part has made a strong case. Intellectual honesty requires stating its limits, because several of them are load-bearing.

It does not prove that customers hate the incumbent. The public satisfaction evidence is three reviews and a middling rating. *[FACT]* The "customers would love to leave" proposition is an **inference from the dissatisfaction-switching gap in an analogous market**, not a documented fact about this one. No bank has publicly left. No specific grievance is documented. *[UNKNOWN]* Any strategy that leads with customer resentment is leading with the least-supported claim available.

It does not prove the platform is technically deficient. The architectural-legacy inference in Chapter 28.3 is built from recruitment postings and a secondary architecture document. *[INFERENCE]* No authoritative technical assessment is public. *[UNKNOWN]*

It does not prove the incumbent is financially distressed in a way that creates an opening. The April 2025 refinancing failed; the July 2025 refinancing succeeded and was oversubscribed. *[FACT]* Sixteen-times leverage constrains flexibility; it does not, on the evidence, prevent the company from funding a product roadmap or defending a market. *[INFERENCE]*

It does not prove the moat is impregnable. The moat protects against replacement. Chapter 30.1 identified the corollary: it does not protect against addition. Whether addition is a real business or a marginal one is the subject of Part Eleven. *[INFERENCE]*

And it does not establish what would actually cause a bank to move. The most important unknown in this part is behavioural: nobody has observed the conditions under which a major agent bank changes its loan servicing

platform, because it has not publicly happened. *[UNKNOWN]* Every strategic argument about displacement in this market is therefore reasoning about an event with no observed instances.

The incumbency case therefore establishes something narrower than "Loan IQ cannot be replaced."

It establishes that replacement is the wrong baseline against which to judge a new capability. The relevant question is whether the capability requires the institution to move the authoritative servicing core, reconstruct its surrounding control environment, or obtain permissions that only the incumbent position currently possesses. *[INFERENCE]*

A capability that can create value without doing those things faces a fundamentally different adoption problem.



PART FIVE

Standards

What three decades of standardization solved, and what it did not

Three trade associations have been working on this market's standards for roughly thirty years. They are competent, well-funded, staffed by practitioners, and supported by every major participant. They have produced model documents, identifier schemes, messaging formats, settlement conventions, and reference-rate transition machinery.

Three decades of standardization changed institutional credit substantially. Documentation became more consistent, identifiers made instruments easier to reference, messaging standards created machine-readable envelopes, and settlement conventions established common expectations.

What standardization did not do was make commercially important agreements uniform, produce one universal identity for every facility, eliminate contextual interpretation, or remove the consent and authority structures embedded in the underlying contract.

The distinction between those achievements and those limits is the subject of this part.

This part asks why. The answer is not incompetence and it is not insufficient effort. It is a structural property of the problem, and articulating that property precisely is the most useful thing this part can do

— because it is the same property that will determine whether any future initiative succeeds.

The documentation project

36.1 What was attempted

The trade bodies publish model documents: credit agreements, trade confirmations, assignment agreements, terms and conditions. The stated theory is that if the market drafts from common starting points, negotiation reduces, comparability increases, and liquidity improves.

36.2 What was achieved

Real and non-trivial things. There are standard trade confirmations for par/near-par and distressed trades. There is a standard form of assignment agreement, which is where the \$3,500 fee lives. [FACT] There are recommended settlement timelines and a delayed-compensation regime. There are joint principles for sustainable lending published across all three bodies. [FACT] There is a joint Administrative Details Form. [FACT] These reduce friction in the mechanical layers of a transaction and they are why the market functions at all.

36.3 What was not achieved

The commercially significant terms. The finding, from practitioner analysis, bears repeating in full because it is the single most important sentence in this part:

The LSTA's model provisions reduce negotiation on boilerplate, but "the provisions that have more commercial effect on the parties (such as mandatory prepayments, business undertakings, financial covenants, representations and warranties, conditions to drawdown, etc.) remain as bespoke to the specific transaction as ever," and the recommended standard documentation forms "are rarely used as a starting draft for negotiation." [FACT]

Direct-lending and syndicated facilities alike "routinely offer bespoke covenant packages, incremental facility baskets and restricted payment provisions negotiated directly with the borrower and sponsor." [FACT] Practitioner drafting commentary documents that even apparently standard clauses are negotiated line by line. [FACT]

Nor did the regional frameworks converge. New York-law and English-law documentation remain distinct; the bodies produced a comparative users' guide and a shared glossary "with the aim of eliminating confusion" — which concedes the persistence of two frameworks. [FACT] Trading English-law loans on New York-law par terms is an active documented legal risk. [FACT] A deliberate convergence attempt — a new LSTA revolving credit facility form — is analysed by practitioners as still notably different from the LMA equivalent. [FACT] And the Asia-Pacific body's documentation uses term-sheet form with **greater** optionality and flexibility tuned to regional practice, meaning less convergence, not more. [FACT]

36.4 Why

The reason is not that the bodies tried and failed. It is that they could not have succeeded, and the mechanism is worth stating precisely.


A credit agreement's economics are the negotiated allocation of risk between a borrower with leverage and lenders with capital. [INFERENCE] A covenant package is a price. A prepayment provision is a price. A margin grid is a price. Standardising them would mean standardising the outcome of a negotiation, which is not something a trade body can do and not something either side would accept. The parts of the agreement that could be standardised are the parts nobody was negotiating over anyway.

This has a direct consequence for software, and it is the central technical fact of the market. **A servicing system must model a facility whose economically significant terms were designed to be unique.** [INFERENCE] Rules-based configuration therefore requires per-deal encoding, performed by a human who has read the agreement, and that encoding is manual,

expensive and brittle. That is the documented reason operations shops fall back to spreadsheets: a spreadsheet is the cheapest available representation of a one-off structure. *[FACT THAT THE FALLBACK OCCURS;*

INFERENCE AS TO THE CAUSE]

The identifier project



If documentation could not be standardised for structural reasons, identifiers should have been easy. They were not, and the failure mode is different and instructive.

37.1 The competing schemes

Four systems coexist. *[FACT]*

- **Loan CUSIPs**, created with the trade association and now "deeply embedded." *[FACT]*
- **LoanX IDs (LXID)**, operated by S&P Global; the identifier franchise reports mapping across more than 400,000 industry identifiers and extended LXIDs to private-credit assets — borrowers, lenders and accounts — in June 2025.

[FACT/VENDOR CLAIM]

- **Bloomberg identifiers**, including the open Financial Instrument Global Identifier standard, which spans asset classes including syndicated loans. *[FACT]*
- **Proprietary internal identifiers** at individual institutions. *[FACT]*

At the entity level the position is worse rather than better. The **Markit Entity Identifier** has been the incumbent since the early 2010s. A **CUSIP Entity Identifier** — a ten-character loan-entity identifier — was subsequently developed by CUSIP Global Services with the trade

association and a bank-owned data platform. *[FACT]* Independent trade press reporting on its launch was headlined, accurately, around the observation that "the path to adoption isn't clear." *[FACT]*

So the market now has two competing entity identifiers for the same purpose, sponsored by different coalitions. *[FACT]*

37.2 The resolution, which is not a resolution

In May 2023 the trade association announced a **Loan Identifier Cross-Reference** product, built and operated with two technology partners, to map identifiers across vendors so that systems "that use different identifiers for loans can communicate reliably." *[FACT]*

Read what that is. It is not a standard. It is a **translation service between competing standards** — an explicit acknowledgment that a single universal identifier was never achieved and that the fix is to translate rather than to unify. *[INFERENCE]*

37.3 Why identifiers failed when they should have succeeded

Identifiers have none of the structural resistance of documentation. Nobody negotiates over an identifier. No party's economics depend on a loan having one number rather than another. It should have been the easy win.

It failed because **an identifier is only valuable if it is the only one, and the parties who could make it the only one each own one.** *[INFERENCE]* The data provider that operates LXID sells data keyed to LXID. The registry that operates CUSIP sells identifier services. Bloomberg distributes an open identifier that supports its own data franchise. Each has a commercial interest in its own scheme's persistence, and none has an interest sufficient to abandon its scheme in favour of a competitor's.

Identifier fragmentation is not a coordination failure. It is a competitive equilibrium. *[INFERENCE]* That distinction matters, because coordination failures can be solved by a convener and competitive equilibria cannot.

The identifier fragmentation is itself a signal of the market's data-infrastructure immaturity, and it is a direct operational cost: every reconciliation across institutions must first resolve which identifier scheme each side is using. *[INFERENCE]*

The messaging project: FpML and the anatomy of a failed standard

This is the market's cleanest documented failure of a published standard, and it is the most instructive episode in this part.

38.1 The attempt

FpML — a financial-products messaging standard stewarded by the derivatives trade association — was extended to loans. The Syndicated Loan Working Group's stated mission was a "common electronic data exchange format" for loan notifications; the loan trade association promoted it; version 5.8 added support for commercial loan processes including loan servicing and position-inventory messages; and message types were defined for loan business events including drawdown, rate set and payments. *[FACT]* The stated intent was to "retire archaic communication methods." *[FACT]*

The industry utility built a messaging service on it: "a safe, secure and automated network for the transmission of industry standard loan messages in FpML format," and it "demonstrated the first end-to-end delivery of FpML messages in the loan industry." *[FACT]* A settlement vendor built a matching product on FpML in 2011. *[FACT]*

38.2 The outcome

The matching product "has yet to win any users." The "lack of full adoption of the FpML protocol prompted platforms to build their own messages instead." Lenders still rely on "fax, telephonic communication and email." "Inertia and budgetary constraints have clearly hampered progress." *[FACT]*

Despite the pilots and the utility hub, "adoption has remained low and manual processes continue to dominate," and the hub still ingested millions of notices by fax and email years after the standard launched. *[FACT]* The messaging portal was acquired by the data provider in 2010 and

absorbed into its notice-data franchise — the standard's principal delivery vehicle became a commercial product of a firm whose business is selling normalised loan data. *[FACT]*

A further signal: the same industry utility later dropped FpML for North American derivatives trade reporting in 2022, in favour of proprietary formats — a different market, but an indication of the standard's contested momentum. *[FACT]*

38.3 The four reasons, and the one that matters

The public record points to four structural reasons rather than a single technical failure. *[FACT]*

One: deals are bespoke. Counterparties span multiple geographies with most deals complex and structured to need, so a single schema struggles to cover them. *[FACT]*

Two: semantic non-standardisation survives syntactic standardisation. Even inside a valid XML envelope, "facility fee" means different things to different senders, and one bank sends start date and end date where another sends start date and tenor. *[FACT]*

Three: adoption was voluntary and partial. *[FACT]*

Four: the network effect ran backwards. A messaging standard is worth nothing to the first adopter and everything to the last. The agent bank that invests in emitting structured messages bears the cost; the lenders who receive them capture the benefit. *[INFERENCE]*

Reason two is the one that generalises, and it is the reason a great deal of subsequent effort has been misdirected. **The market's problem was never that its data lacked a format. It was that its data lacked a shared meaning.**

[INFERENCE] A syntactic standard cannot supply semantics. This is why the extraction problem in Part Nine is a semantic problem, why "we can parse the PDF" is not a solution, and why the interesting question about AI in

this market is not whether it can read a document but whether it can resolve what the document means in a way that another institution will accept. *[INFERENCE]*

Reason four is the one that determines strategy. **A standard whose costs and benefits fall on different parties will not be adopted, regardless of merit, unless somebody compensates the cost-bearer.** *[INFERENCE]* Chapter 41 generalises this.

The settlement project: twenty-five years

The trade bodies have attacked settlement time for a quarter-century.

The record:

- **Delayed compensation** introduced in 1998 to make late settlement economically neutral. *[FACT]*
- **1 September 2016:** the regime shifted from no-fault to requirements-based, because the association concluded the no-fault design was counterproductive — "Buyers may not be motivated to settle trades when they are guaranteed delayed compensation." *[FACT]*
- **Buy-in/sell-out provisions** at T+20, criticised in the same analysis as "too long after a trade date to mitigate delayed settlement risks in less dramatic situations." *[FACT]*
- **July 2023:** distressed provisions folded into the par standard terms because participants were trading distressed loans on par documents. *[FACT]*
- **May 2026:** late-payment fees introduced for loan trade transactions. *[FACT]*
- **2026:** an active reform programme on standardisation, automation, and credit-agreement provisions to shorten settlement, targeting consent tracking, agreement updates and lender communications — with distressed settlement singled out and no fixed settlement mandate. *[FACT]*

The measured outcome across that period: twenty-three per cent of par trades settling within T+7 in 2016 *[FACT]*; approximately twenty-nine per cent within T+7 and twenty-seven per cent beyond T+20 on 2021 data *[FACT – LAW-FIRM SOURCED]*; an average of roughly 19.5 days in 2023, seven per cent better than 2021–22 *[FACT]*. Current figures are member-gated. *[UNKNOWN]*

Twenty-five years of sustained effort by competent institutions produced an improvement measured in single-digit percentages against a target the market itself set. *[INFERENCE]*

39.1 The association's own diagnosis

The most important document in this chapter is a trade-association publication whose title is its argument: **"Changing credit agreement provisions is necessary to shorten settlement times."** *[FACT]*

That is the standard-setter conceding that the bottleneck is not in the systems, the platforms or the conventions it controls, but in the **contracts** — the consent regimes, the disqualified-lender machinery, the transfer restrictions negotiated deal by deal between borrowers and lenders who have their own reasons for wanting them. *[INFERENCE]*

The structural fix has been repeatedly urged and never achieved: electronic platforms, universal CUSIPs, standardised disqualified-lender lists. *[FACT]*
Central electronic settlement of the kind that exists in securities markets remains absent. *[FACT]*

The one that worked

Among the coordinated changes surveyed in this paper, one succeeded at scale, and it is the exception that establishes the rule.

The transition from LIBOR to risk-free reference rates required every floating-rate loan agreement in the market to be amended or to fall back to a new rate. The trade association co-chaired the relevant business-loans working group of the official-sector-convened rate committee alongside a banking association and a law firm. Recommended fallback language for US dollar business loans was published on 25 April 2019, offering an "amendment" approach and a "hardwired" approach. By December 2021 the association had a full suite of concept documents for the successor rate.

[FACT] European coordination proceeded in parallel through changes to secondary trading documentation. [FACT]

It worked. The market transitioned.

Why it worked is the whole point. The transition was forced by an external deadline set by a party outside the market: a regulator announced that the reference rate would cease to be published. [FACT] Non-adoption was not an option, because the alternative was a contract referencing a rate that did not exist.

Set that against every other initiative in this part. Documentation convergence: optional. Identifier unification: optional. FpML adoption: optional. Faster settlement: optional. In each case a participant could decline and continue operating normally, and in each case a meaningful number did.

The market's trade bodies can coordinate the market when an external authority removes the option of not coordinating. They cannot coordinate it otherwise. [INFERENCE]

FIGURE 5

The one coordination that worked, and what it required



The reference-rate transition shows that this market can coordinate when an external authority removes the option of not coordinating — and records what that took: a decade, a convened governance structure, common conventions and a legislative backstop. Chapter 40 treats it as the counterexample that establishes the rule. Scale figures are official-sector and trade-body estimates. [FACT]

A theory of standardization failure in this market

INITIATIVE	COST FALLS ON	BENEFIT FALLS ON	RESULT
FpML	Agent	Lenders	Not adopted
Faster settlement	Agent	Counterparties	Three-week settlement
Universal identifiers	Scheme owners	Everyone else	Cross-reference service
Transparency	Dealers	Buy side	No tape
LIBOR transition	All	All	Adopted (external deadline)

The episodes in this part are usually narrated as a list of disappointments. They are better understood as instances of one mechanism, and stating the mechanism precisely gives it predictive value.

41.1 The mechanism

Proposition: in this market, a standardization initiative fails when its costs fall on the agent and its benefits fall on the lenders, unless the cost-bearer is compensated or compelled.

This is offered as a mechanism to be tested, not a law. Test it against each case.

FpML. The agent bank must invest to emit structured messages. The lenders receive them and save re-keying. The agent's benefit is minimal — it already knows what it sent. **Cost: agent. Benefit: lenders. Result: not adopted.** *[INFERENCE]*

Faster settlement. The agent must staff to review and acknowledge trades within a business day. The benefit accrues to buyers and sellers whose capital is tied up. The agent is not paid more for being fast, and the assignment fee is the same either way. **Cost: agent. Benefit: counterparties.**

Result: three-week settlement. [INFERENCE]

Universal identifiers. Adoption requires the data vendors, each of whom owns a scheme, to subordinate their own. **Cost: scheme owners. Benefit: everyone else. Result: a cross-reference service.** [INFERENCE]

Transparency. A public tape transfers value from dealers to investors, and the empirical evidence from the corporate bond market says the transfer is substantial. [FACT] The dealers control the market's infrastructure decisions. **Cost: dealers. Benefit: buy side. Result: no tape.** [INFERENCE]

LIBOR transition. The cost fell on everyone and the alternative was contractual chaos for everyone. **Cost: all. Benefit: all. External deadline. Result: adopted.** [INFERENCE]

The pattern holds across all five. That is not proof. Five observations, selected by the author from a population the author also defined, is a weak test, and the mechanism was formulated after the cases were known. What can be said is that the mechanism is independently plausible on incentive grounds, that it explains all five cases without special pleading, and that it makes falsifiable predictions — which are set out in §41.3 precisely so that a reader can check them against events the paper did not use to build the theory. [INFERENCE]

41.2 The corollaries

Three follow, and they should be applied to every proposal about this market — including the industry agenda in Part Twelve and the entrant paths in Part Thirteen.

Corollary 1: Any initiative requiring the agent to bear cost for others' benefit will fail, absent compensation or compulsion. *[INFERENCE]* This kills industry-wide messaging standards, voluntary transparency, and most settlement reform. It is why the association's own 2026 programme is aimed at credit-agreement provisions rather than at agent behaviour.

Corollary 2: Initiatives that route around the agent, or that pay the agent, can succeed. *[INFERENCE]* The bank-owned data platform pays for agent data with equity and control — the agent banks are the founders. That is compensation, structured as ownership. It is the reason that model has traction where FpML did not. *[INFERENCE]*

Corollary 3: One-customer-at-a-time approaches escape the mechanism entirely. *[INFERENCE]* A product sold to a single institution, delivering value to that institution, requires no coordination, no adoption threshold, and no cost-benefit alignment across parties. This is the strongest structural argument available for a coexistence strategy, and Part Eleven leans on it heavily; Chapter 91 asks what it implies for an institution acting alone. It is also the argument's principal limitation: a product that requires no coordination also generates no network effect, and therefore no defensibility from network effects. *[INFERENCE]*

This third corollary deserves particular attention because it changes the sequencing of modernization.

The first two corollaries describe how a market can coordinate. The third identifies a class of improvements that do not require market coordination at all.

That distinction will become materially more important once Part Nine introduces a capability that an institution can apply to documents and information it already possesses. *[INFERENCE]*

41.3 What this predicts


If the mechanism is right, then:

- The 2026 settlement reform will produce marginal improvement, because it targets credit-agreement provisions negotiated between borrowers and lenders rather than agent behaviour, and borrowers have reasons to want transfer restrictions. *[HYPOTHESIS]*
- Identifier fragmentation will persist and be managed through translation rather than resolved. *[HYPOTHESIS]*
- No public loan-price tape will emerge absent regulatory compulsion. *[HYPOTHESIS]*
- The bank-owned data utility will continue to gain coverage, because it compensates the cost-bearer with ownership, and its progress will be gated by how many agents accept the bargain rather than by its technology. *[HYPOTHESIS]*

Each of these is falsifiable, and a reader who expects a different outcome has a specific place to disagree.

FIGURE 6

Where the cost falls, and where the benefit does



Coordination costs in this market concentrate on a few institutions while the benefits distribute across many — the mechanism Chapter 41 states as its first corollary. The initiatives shown are those examined in Parts Six and Seven, with the reference-rate transition included as the case where compulsion replaced alignment. [INFERENCE]

VI

PART SIX

How technology evolved around the incumbents

The incumbent servicing platforms were installed by the late 1990s and have not been displaced. Everything the market has built since has been built *around* them. This part traces six waves of that building, in order, and closes by extracting the pattern they share.

The pattern is the useful part. Each wave was staffed by capable people with real capital and a genuine problem, and each produced a partial result of a specific and repeating shape. Knowing that shape is the closest thing this paper can offer to a predictive tool.

Wave one: workflow digitization

The first response, running roughly from the late 1990s through the 2000s, was to digitise individual workflows without touching the core.

Deal sites replaced physical data rooms for primary syndication. Document distribution platforms replaced couriers. Notice management moved from paper files to document repositories. The incumbent servicing vendors themselves shipped syndication and deal-management products alongside their servicing cores. *[FACT]*

The archetype of the wave is a notice-management product: "an online document management service that processes, categorizes and stores agent notices related to bank loan holdings," letting loan managers "handle inbound agent notices." *[VENDOR CLAIM]*

Read that description carefully. It is a **capture-and-sort layer over an unstructured inflow, not a fix to the source.** *[INFERENCE]* The notices still arrive by fax and email, still carry inconsistent semantics, and still require human interpretation. The product makes them easier to file.

That is the signature of wave one and it recurs: **the workflow around the problem was digitised; the problem was not.** *[INFERENCE]* The market got better filing cabinets.

Wave two: data and pricing consolidation

The second wave, running through the 2000s and 2010s, was commercial rather than technological: if the market's data was fragmented and hard to obtain, someone could aggregate it and sell it back.

The consolidation chain is documented in Chapter 12.1 and does not need repeating. Its outcome: pricing, reference data, identifiers, settlement and administration under a single owner, with regulatory intervention at the index and commentary layer only. *[FACT]*

The products this wave produced are real infrastructure. Daily mark-to-market pricing on thousands of leveraged loan facilities sourced from sixty-plus sell-side desks. *[VENDOR CLAIM]* Loan reference data covering corporate actions and terms. Identifier services. A notice-data hub receiving seven million notices a year. A loan reconciliation product automating lender-agent position reconciliation. *[VENDOR CLAIM]* Competing pricing services from a terminal provider and an exchange group. *[FACT]*

Two consequences.

The wave solved the data-availability problem and did not solve the data-authority problem. *[INFERENCE]* A lender can now buy a price, an identifier and a normalised notice feed. It still cannot verify that its position matches the agent's without asking the agent.

It concentrated the market's data leverage. The firm that aggregates the notices, operates the settlement platform, sells the reference data and administers the CLO books occupies a position from which a great deal of the market's operational reality is visible. *[INFERENCE]* Any new entrant in loan data is entering a market with a very large incumbent that has, by acquisition, already occupied most of the available positions.

Wave three: settlement and reconciliation utilities

The third wave attempted the structural fix: build shared utilities for settlement and reconciliation.

Settlement. ClearPar became the dominant electronic settlement ecosystem, described in a securities filing as the primary platform for electronic confirmation, documentation and settlement of syndicated loans in the United States, settling "substantially all" leveraged syndicated loans and supporting the European market. *[FACT]* A competing platform, founded in 2005 by former bank executives, holds a smaller share. *[FACT]* Trades not on a platform follow paper requirements. *[FACT]*


The result: settlement is electronic in the sense that documents and funds route through a platform, and settlement still takes three weeks, because the platform routes a process whose bottleneck is consent and review (Chapter 22.3). *[INFERENCE]* **Electronifying a workflow does not accelerate a workflow whose latency is human decision.** *[INFERENCE]*

Reconciliation. An industry-owned utility launched a position-reconciliation service in 2008 and a messaging service on the loan messaging standard, reaching roughly 2,300 funds and, on some accounts, roughly 400 asset managers and 6,000 funds. *[FACT]* Its rationale was explicit: before it, "lenders and banks reconciled loan positions independently of one another, a process prone to errors, miscalculations and delays." *[FACT]* Its messaging portal was acquired by the commercial data provider in 2010 and absorbed. *[FACT]*

The utility survives and provides real value. It also demonstrates the wave's limit. A daily position reconciliation tells a lender that its record differs from the agent's. It does not tell the lender why, does not resolve the difference, and does not prevent the next one. *[INFERENCE]* The break is detected rather than eliminated, and detection is the cheap half.

Wave four: distributed ledger, and the graveyard

THE VENTURE'S ARC



Between roughly 2016 and 2019 the market convinced itself that the answer was a shared ledger. The reasoning was superficially compelling: if the problem is that many parties each maintain their own version of the same facts, then give them one version they all maintain together.

The wave failed. Its failure is the most instructive episode in the market's technology history and deserves detail.

45.1 The syndicated-loan venture

A distributed-ledger company built an enterprise smart-contract platform aimed at the syndicated-loan back office and formed a joint venture with a capital-markets data firm to move the secondary syndicated-loan market onto a shared ledger. *[FACT]*

Then the market consolidated underneath it. The data firm was acquired in 2018 by the same group that had by then assembled the market's loan-data infrastructure. Backers pulled support; the venture stalled; its chief executive resigned in November 2019; and a Delaware court ordered it wound down. *[FACT]*

The DLT company sued and **won**. The Delaware Court of Chancery ruled that the acquisition breached a non-competition agreement. Coverage cites a headline award of approximately \$70 million to the company's half share, with estimates running to roughly \$140 million at the joint-venture level; a settlement of **\$53 million** was paid at the end of 2021. *[FACT – THE EXACT FIGURE VARIES ACROSS SOURCES AND SHOULD BE TREATED AS UNSETTLED]*

The company then **filed for Chapter 11 on 1 December 2022**. Legal costs were "massive," there was "no cash left," and a \$2.3 million secured loan came due. *[FACT]* Its sole secured creditor acquired the assets via a stalking-horse bid of roughly \$2.6 million, approved in June 2023. *[FACT]*

A technically capable venture, in the right market, with a demonstrated product, won its litigation and died anyway – killed by partner-consolidation risk, litigation dependency and cash burn, not by any technical failure. *[FACT]* Chapter 52 draws the strategic lesson.

45.2 The incumbent's ledger

The dominant servicing vendor built its own DLT product on an enterprise ledger platform, letting agent banks publish real-time credit-agreement, accrual and position data to lenders, fed from the servicing core. It was the first application to go live on that ledger platform. *[FACT]* Adoption reached roughly seven pilot banks, characterised as around twenty-five per cent of global syndicated-loan agent volume, with further sign-ups through 2018 and 2019. *[FACT/VENDOR CLAIM]*

No public reporting on its blockchain deployment exists after approximately 2019. Whether it still runs on the original ledger, was folded into the vendor's interface layer, or was quietly retired is not verifiable from public sources. *[UNKNOWN]* The silence is itself notable.

What is documented is that in October 2021 the product was integrated with a data-collaboration network's data layer, so that agents publish contract- and transaction-level data and lenders reconcile and monitor it, with updates flowing back to the servicing core. *[FACT]* The ledger product's

afterlife, in other words, is as a data source for a conventional coexistence network — which is a fair summary of what the whole wave amounted to. *[INFERENCE]*

45.3 The verdict from outside

The independent evidence on why DLT failed in loans is unusually clear.

Only **17.7 per cent** of one trade body's members in a November 2020 survey were using or considering blockchain, against **60.4 per cent** for plain electronic document and transaction platforms. *[FACT]* Enthusiasm never converted.

The named barriers: legal and regulatory uncertainty, in that existing law does not make DLT "legally safe" for loans; network interoperability and fragmentation; scalability; and — the finding that matters most — that **for private permissioned networks a centralized system is often faster, cheaper and better.** *[FACT]*

The academic and industry consensus is that **the binding constraint was coordination, legal enablement and network effects among agent banks — not compute.** *[FACT]*

Even the healthiest adjacent DLT venture, a collateral-mobility platform, is migrating off the original ledger platform to a different network — a signal about that platform's enterprise trajectory. *[FACT]*

45.4 What the wave got wrong

The diagnosis is worth stating precisely because a version of it recurs in every subsequent wave.

DLT answered the question "**how do mutually distrustful parties agree on a shared state?**" That is a real question and DLT is a real answer to it.

But it is not this market's question. *[INFERENCE]* In a syndicated loan there is no dispute about who decides. The agent decides. The register is authoritative by contract. The parties are not mutually distrustful peers

requiring consensus; they are a principal and a set of counterparties in a hierarchy of authority established by a negotiated agreement.

What the market actually needs is not agreement on shared state. It is **timely, verifiable propagation of authoritative state from the party that holds it to the parties that need it, with enough provenance that the recipients can rely on it.** *[INFERENCE]* That is a distribution and evidence problem, not a consensus problem, and Chapter 71 develops the distinction at length.

The wave failed because it solved the wrong problem with impressive technology. *[INFERENCE]*

Wave five: the bank consortia

The fifth wave, from 2022, is the market's most successful modernization effort to date and the most important competitive fact for anyone considering entry.

Its innovation is not technological. It is **ownership**.

46.1 The data platform

Versana was announced on 16 March 2022 by four founding banks — JPMorgan, Bank of America, Citi and Credit Suisse — to digitise what was then framed as the "\$5 trillion" syndicated loan market. *[FACT]* It captures administrative agents' books-and-records loan data in real time via application programming interfaces and centralises it for lenders, attacking fragmentation and the twenty-day settlement cycle. *[FACT]* The platform went live in December 2022. *[FACT]*

Backing broadened through 2023 to include Deutsche Bank, Morgan Stanley, US Bancorp and Wells Fargo. *[FACT]* Coverage grew: more than 1,500 facilities and roughly \$900 billion of commitments by October 2023, against a stated target of roughly ninety per cent of broadly syndicated loans and about 6,000 facilities by end-2024. *[FACT/VENDOR CLAIM]* JPMorgan achieved a straight-through-processing milestone in June 2024. *[FACT]* Reported active facility coverage subsequently exceeded **\$4.1 trillion** in notional value, with the market framing having grown from "\$5 trillion" in 2022 to a "\$9 trillion broadly syndicated loan and private credit" market by 2026.

[VENDOR CLAIM]

Buy-side monetisation began in January 2025 with the first buy-side subscriber, and a reconciliation module launched the same month to electronically match lender and agent records and display the transaction notices causing differences. *[FACT/VENDOR CLAIM]*

The April 2026 funding round is the single most important competitive development in this market. Versana closed \$43 million on 30 April 2026, led by BNP Paribas, with new investors including **Apollo**, a ratings agency's venture arm, an insurer's venture arm and a financial-technology investor, plus follow-ons from every existing bank owner. Total raised exceeded \$125 million. *[FACT]*

The significance is not the money. It is that a bank-owned consortium has added a **mega-manager** to its cap table and is explicitly courting the buy side and private credit. *[FACT]* Chapter 78 develops what that means.

46.2 The trading venue

Octaura was announced on 14 June 2022, backed by seven dealer banks and a ratings and analytics firm, to build an open electronic trading platform for syndicated loans and CLOs; originally built by two of the banks. *[FACT]* It reached a proof point before it reached a product: **the market's first fully electronic syndicated loan trades were completed in February 2023**, ahead of the **loan trading platform's launch in April 2023**. *[FACT]* Octaura announced the closing of a \$46.5 million funding round on 9 June 2025, adding further bank and institutional investors; CLO trading launched September 2025 with three protocols, twelve dealers and more than fifty buy-side firms. *[FACT]*

That sequence — a live trade demonstrated two months before the venue opened — is a small but revealing detail about how this market absorbs new infrastructure. The technical capability was not in doubt and was not the gating item; what took the intervening period was the surrounding apparatus of participants, documentation and operational readiness.

[INFERENCE] It is the same ordering the rest of Part Six describes at larger scale.

Traction: dealer network from three to roughly twenty-five to thirty and buy-side from thirty-four to more than 170 between 2023 and 2026; secondary loan market share from roughly one per cent in the first quarter

of 2024 to approximately 4.6 per cent in the first quarter of 2025 and roughly 7.7 per cent by January 2026; roughly \$1 billion a week traded since early 2025; a first quarter of 2026 at \$19.6 billion. *[VENDOR CLAIM – SELF-REPORTED, DIRECTIONALLY CORROBORATED IN TRADE PRESS]*

These are among the few hard adoption numbers any utility in this market has disclosed, and they show real progress from a very low base. *[INFERENCE]*

46.3 The coexistence network

AccessFintech is the third model and the one this paper regards as most consequential for strategy.

Founded in 2016 by a veteran of post-trade matching infrastructure. *[FACT]* Its Synergy network is a cloud-native, multi-party data-collaboration layer: participating institutions publish transaction, position, cash and reference data into a shared network that normalises it into a common language, then surfaces the same data and exceptions to all counterparties so they investigate and resolve collaboratively – **without replacing existing source systems.** *[FACT]*

Funding: approximately \$97 million across four rounds, with a \$20 million Series B in October 2020 led by a venture firm with Deutsche Bank joining alongside JPMorgan, Citi and Goldman Sachs; and a \$60 million Series C in September 2022 led by a growth investor, with BNY Mellon and Bank of America new and BNP Paribas joining separately. *[FACT]* No Series D appears on the public record. *[UNKNOWN]*

Scale: more than 250 institutions across buy side, sell side, asset servicers and order-management platforms; more than one billion transactions contributed per month. *[VENDOR CLAIM]* Bank-reported outcomes of approximately thirty per cent fewer trade fails and seventy-six per cent less email traffic after a year, attributed to two major banks. *[FACT – REPORTED IN INDEPENDENT TRADE PRESS]*

Its loans coverage is real and specific. The network is organised around securities, private markets including loans, and derivatives. In October 2021 it integrated with Finastra's Fusion LenderComm, so that loan agents publish contract- and transaction-level data and lenders reconcile and monitor it, with updates flowing back into Loan IQ. *[FACT]* It operates a dedicated loan data management service premised on syndicated loans' "lack of transparency around agent data" where "absence of data reconciliation leads to incorrect payments and rework." *[VENDOR CLAIM]* It has a partnership with a corporate trustee to streamline private-credit lifecycle management. *[VENDOR CLAIM]*

In November 2025 it announced a strategic partnership with a major buy-side portfolio platform to connect that platform's global community to the network. *[FACT]*

46.4 Three bets, compared

	BANK DATA PLATFORM	TRADING VENUE	COEXISTENCE NETWORK
Bet	Go to the agent's source system; bypass notices	Electronify execution	Sit alongside every system; reconcile across counterparties
Relationship to incumbent core	Reads from it	Orthogonal	Reads from and writes back to it
Ownership	Bank consortium plus a manager	Bank consortium	Bank-and-VC backed
Adoption metric	Facility coverage and notional	Share of secondary volume	Institutions and transactions
Structural strength	Compensates agents with ownership	Neutral venue with dealer protocols	Requires no one to change anything
Structural weakness	Progress gated by agent onboarding	Preserves dealer intermediation	Detects and normalises; does not read contracts or take over

The comparison is the strategic map of the current market. *[INFERENCE]* Note in particular the last row of the final column: the coexistence position is occupied, by a firm with a decade of network effects, the bank cap table, the incumbent integration and a buy-side distribution partnership — and its identified limitation is that it reconciles data rather than reading the

underlying agreements or taking over any function. Whether that limitation is a genuine opening or a description of the natural boundary of the position is the question of Part Eleven. *[INFERENCE]*

Wave six: artificial intelligence and document extraction

The sixth wave, from roughly 2023, applies large language models to the market's Category 2 problem: reading the bespoke agreement.

The field is crowded. Entrants include deep-research platforms extracting loan terms and covenants from credit documents; unstructured-document processors with human-in-the-loop verification; clause-extraction tools now inside legal-technology suites; alternative-asset document ingestion platforms; AI-native loan origination systems for commercial banks; document readers for credit agreements built into securitisation analytics; valuation-first extraction platforms for private credit; and general legal and research AI applied to credit documents. *[FACT]*

The consolidation is already under way and is the most telling feature of the wave. **Finley**, a debt-capital-management platform, was acquired by Concord in February 2026. **Solvas**, Deloitte's CLO accounting and compliance suite, was acquired by Alter Domus. **Lumonic**, an AI-native private-credit monitoring platform, was acquired by Morningstar in March 2025 and positioned as "the only AI-native private-credit spreading, end-to-end, with source-cell traceability." **Kira Systems**, the clause-extraction pioneer, sits inside Litera. **Evisort**, a contract-intelligence platform, was acquired by Workday. *[FACT]*

Funding continues: **73 Strings**, a valuation-and-extraction platform, raised \$55 million in February 2025; **Casca**, an AI-native origination platform, raised \$29 million in 2025; **Setpoint**, a capital-markets operating system for asset-backed lending, reports more than \$50 billion of financed assets a year with backing from Citi Ventures. *[FACT]*

Two conclusions.

Extraction is becoming table stakes rather than a moat. *[INFERENCE]* The evidence is the consolidation rate: five acquisitions of extraction-adjacent companies since 2021, three of them inside the last eighteen months, including one by the market's own data incumbent. When incumbents buy a capability rather than compete with it, the capability is being absorbed into the baseline. *[INFERENCE]* Note especially that the acquired private-credit monitoring platform is positioned on **source-cell traceability** — that is, on provenance, which is precisely the differentiator a coexistence thesis would claim. A data incumbent already owns it. *[FACT]*

One entrant is doing something structurally different and deserves separate attention. Setpoint performs **collateral document verification, borrowing-base computation and compliance monitoring** — that is, it verifies the underlying assets rather than merely extracting terms from a document. *[VENDOR CLAIM]* That is a closer competitor to any verification-oriented thesis than the extraction field generally, and Part Thirteen treats it as such. *[INFERENCE]*

The pattern across waves

1990s–2000s	2008–10	2008 on	2016–19	2022 on	2023 on
Workflow digitization	Data & pricing consolidation	Settlement & reconciliation utilities	Distributed ledger — failed	Bank consortia — most successful	AI extraction — consolidating fast

WAVE	WHAT IT REQUIRED, AND WHAT HAPPENED
1 • Workflow digitization	Required nothing of anyone else. Succeeded — as a filing cabinet, not a fix.
2 • Data consolidation	A lender could buy a price without anyone's cooperation. Succeeded; solved availability, not authority.
3 • Settlement / reconciliation	Routed documents through a platform without adopting anything. Succeeded at detection, not resolution.
4 • Distributed ledger	Required a network of banks to write to shared state. Failed — critical mass never formed.
5 • Bank consortia	Compensated the cost-bearer with ownership. The market's most successful modernization.
6 • AI extraction	Required nothing of anyone else. Succeeding — and consolidating into incumbents.

Six waves. One shape.

Across the six waves surveyed here, those that succeeded did so by adding a layer that required no other participant to change behaviour. Data aggregation succeeded because a lender could buy a price without anyone's cooperation. Settlement platforms succeeded because the parties to a trade could route documents through one without the market adopting anything. The coexistence network succeeded because it explicitly does not replace source systems. *[INFERENCE]*

Those that failed required coordinated adoption. The messaging standard required agents to emit and lenders to consume. The distributed ledger required a network of banks to write to a shared state. Both

were technically sound and both required a critical mass that never formed. *[INFERENCE]*

None of the six displaced the core. After three decades of building, the agent's servicing platform still holds the register, still generates much of the operative servicing state and its notices, and still determines the facts for which the credit agreement assigns the agent authority. Not one of the six waves attacked it, and the one venture that came closest — the joint venture aiming at the secondary market — died. *[INFERENCE]*

Each wave addressed the symptom that the previous wave's success made visible. Digitised workflows made the data problem visible; data aggregation made the authority problem visible; settlement platforms made the consent problem visible; the ledger wave made the coordination problem visible; the consortia made the ownership problem visible; and the AI wave is making the semantic problem visible. *[INFERENCE]* The market has been peeling an onion, and the layers have got harder rather than easier.

And the residue is constant. Across six waves and three decades, three things have not changed: the agent holds the register; the agreement is bespoke; and settling requires consent. *[FACT]* Those are the Category 3 and Category 5 problems of Chapter 26, and no wave has touched them.

The honest summary is uncomfortable. **On the record assembled here, the market adopts what does not require it to change, and declines what does.**

[INFERENCE] Any thesis about this market must be evaluated against that finding, and a thesis that requires coordinated change is, on three decades of evidence, a thesis whose merit is not the binding constraint on its adoption.

FIGURE 7

Six waves of modernization, and what survived them

A more connected market; the same underlying problem. Each wave solved real problems and added lasting capability, and the institutional architecture underneath has largely survived.

1

1990S – EARLY 2000S

Systematize the loan — core loan platforms

Put complex lending into purpose-built systems.

Transformational · established the foundation that still exists today

Core loan servicing platforms — the category the market consolidated around, and within it Loan IQ and ACBS (Chapters 27–29)

2

2000S – 2010S

Digitize the workflow — electronic processes and market infrastructure

Move from paper to electronic. Add settlement, notices, reconciliation and data tools.

High impact · materially improved efficiency and reduced operational risk

Electronic settlement, notice-distribution and reconciliation platforms, and market-data providers — including the settlement platform examined in Chapter 39

3

2010S

Standardize and connect — standards, data and integration

Make systems speak the same language. Expand connectivity across participants.

Meaningful · improved connectivity, but did not standardize meaning, authority or state

Standards and messaging initiatives — the documentation, identifier and messaging projects of Part Five, including FpML (Chapter 38)

FIGURE 7, CONTINUED

4	2016 – 2020 Create shared state — distributed ledger and consortium models Explore new models for shared data and real-time information exchange. Limited / mixed · proved concepts and increased transparency, but could not achieve broad participation <i>Distributed-ledger and consortium initiatives — the wave surveyed in Chapter 45, including Symbiont (Chapter 52)</i>
5	2022 – PRESENT Centralize authoritative data — agent-sourced data platforms Provide trusted, real-time data from authoritative sources such as administrative agents. High impact, growing · solves important data-access problems, but does not resolve interpretation, consent or institutional incentives <i>Agent-sourced data platforms — including Versana and its founding agent banks (Chapter 46)</i>
6	2024 – PRESENT Understand the loan — AI and semantic interpretation Turn unstructured agreements and documents into usable institutional data; automate extraction, comparison and analysis. Emerging · significant potential; the real test is whether it changes institutional constraints <i>AI and document-interpretation systems — incumbent platforms with AI intake, general-purpose model providers, and document- and contract-analysis entrants (categories, not a vendor census)</i>

THE PERSISTENT FOUNDATION

Bespoke contracts Every deal is different. Language, terms and exceptions vary.	Distributed authority Multiple parties. No single owner.	Institution-specific state Each institution maintains its own version.
Consent Change requires many participants.	Coordination Complex, multi-party processes.	Economic incentives Misaligned incentives slow standardization.

The industry did not fail to invest in modernization. It accumulated thirty years of it. The machinery changed; the institutional architecture largely survived.

Wave dating, verdicts and participant positioning are illustrative and follow the survey in Parts Four through Six; naming a platform is not an assessment of it.

Each wave solved real problems and left lasting capability; none displaced the servicing core or removed the constraints beneath it. Chapter 48 sets out the pattern the waves share. Platforms and vendors named are those discussed in Parts Four through Six; wave positioning is illustrative. [INFERENCE]

VII

PART SEVEN

Why collaboration repeatedly failed

The institutional credit market has a peculiar relationship with cooperation. Its largest participants compete ferociously for every mandate and jointly own most of its shared infrastructure. They fund each other's utilities, sit on each other's boards, and appear side by side in the founder list of nearly every market-infrastructure venture of the last decade — while simultaneously fighting each other for the same deals.

This is not hypocrisy. It is a coherent strategy with a documented logic, and understanding it is a prerequisite for any assessment of entry into this market. The strategy is highly effective at a particular class of problem and structurally incapable of solving another class. This part identifies which is which.

The consortium as institutional form

The bank-owned consortium is one of the market's most reliable institutional responses, and it long predates the loan-technology consortia.

The precedents are consistent in form. A secure messaging venture founded in October 2014 by fifteen firms including the major dealer banks and a large asset manager, with \$66 million of funding. A payments network owned by seven of the largest US banks that moved more than \$1 trillion in 2024. A clearing house owned by roughly twenty-one or twenty-two of the largest banks, whose four biggest members are jointly building a tokenised-deposit network for a 2027 target. *[FACT]*

The loan-market consortia follow the same template. The data platform's founders and later investors include JPMorgan, Bank of America, Citi, Credit Suisse, Deutsche Bank, Morgan Stanley, US Bancorp, Wells Fargo, BNP Paribas and Barclays. The trading venue's backers include Bank of America, Citi, Credit Suisse, Goldman Sachs, JPMorgan, Morgan Stanley, Wells Fargo, Barclays, Deutsche Bank, BNP Paribas and a ratings and analytics firm. The coexistence network's investors include JPMorgan, Citi, Goldman Sachs, BNY Mellon, Bank of America, Deutsche Bank and BNP Paribas. *[FACT]*

The same names appear in all three. *[FACT]*

49.1 The stated logic

The strategy has been articulated explicitly in independent analysis. Large banks repeatedly form cooperative — "you could argue, anti-competitive" — utilities to neutralise threats to their franchises rather than let an independent vendor capture the space. The consortium product "doesn't need to be profitable; it just needs to preserve or strengthen the banks' existing competitive moats," and it is "usually good enough, but not best-in-class." *[FACT — THIS IS INDEPENDENT ANALYSIS, NOT THE BANKS' OWN CHARACTERISATION]*

That is a remarkably clear statement of intent, and its two halves both matter.

"Doesn't need to be profitable" means a consortium competitor operates without the unit economics constraint that binds a venture-backed entrant. It can price at cost or below indefinitely, because its return is measured in preserved franchise value rather than in enterprise value. *[INFERENCE]*

"Good enough, but not best-in-class" means the consortium's product quality ceiling is genuinely lower — which is the entrant's only structural advantage, and a thin one. *[INFERENCE]*

49.2 What consortia are good at

The pattern across the precedents is that consortia succeed at problems with three properties: the problem is **shared infrastructure** rather than competitive differentiation; the **benefit is symmetric** across participants; and **no participant loses** from the solution. *[INFERENCE]*

Messaging, payment rails, clearing, and loan data all fit. Every bank benefits from cheaper, faster, more reliable plumbing, and none loses market share to another because of it.

49.3 What consortia cannot do

They cannot solve a problem where the members' interests diverge. *[INFERENCE]*

The clearest instance is transparency. A public price tape would transfer value from dealers to investors (Chapter 21.3, Chapter 53). The consortium members are the dealers. A consortium of dealers will not build a tape that costs dealers money, and none has. *[INFERENCE]*

The second instance is anything that reallocates value among the members themselves. Consortia pool infrastructure; they do not resolve distributional questions. *[INFERENCE]*

The third — and this is the one that matters most for entry — is that a consortium is a **defensive** institution. It exists to prevent a chokepoint from being captured. Its success condition is that no outsider controls the layer, not that the layer works optimally. *[INFERENCE]* That is why the analysis quoted above is right that the product need not be best-in-class: being adequate and owned is the objective.

Coopetition: the documented rule

The market's operating rule can be stated in one line and is supported across a decade of behaviour:

Where the problem is shared infrastructure — data, settlement, messaging, connectivity — the top banks pool it into a jointly owned utility so that no single rival captures a chokepoint. Where the problem is winning the client mandate, they compete ferociously. *[INFERENCE, FROM THE DOCUMENTED MEMBERSHIP PATTERNS]*

The pattern is observable rather than asserted. The same small set of institutions recurs across the market's cooperative structures — the founding groups of the agent-data platform, the electronic trading venue, the settlement and reconciliation utilities, and the trade association itself are drawn from substantially overlapping membership (Chapters 44, 46, 49). *[FACT]* The same institutions compete without qualification for the mandates those structures serve. Coopetition here is not an attitude; it is a division of the map, and the map is stable enough across a decade to be treated as a rule rather than a tendency. *[INFERENCE]*

The boundary is where a new entrant lives or dies, and it is not obvious in advance. A product that looks like a bank-specific workflow tool may be reclassified as shared infrastructure the moment it becomes valuable; a product that looks like infrastructure may be tolerated as a niche indefinitely. Three questions locate a proposition on the map before the market does, and they are worth asking early because the answer determines the competitive response rather than the product roadmap. *[INFERENCE]*

Does it sit between the banks, or inside one of them? Anything that must be used by several institutions to be useful is, by construction, shared infrastructure. Anything whose value is complete inside a single institution is not — and this is the one property an entrant can actually choose.


Would one bank owning it disadvantage the others? This is the test the banks themselves appear to apply. A chokepoint that a rival could own gets pooled; a capability that confers no leverage does not.

Does it touch the client relationship? Origination, pricing and the mandate are competitive territory and are not pooled at any level of value.

The practical implication for an entrant is that **success is the trigger**.

[INFERENCE] A vendor that fails to gain traction is ignored. A vendor that succeeds in a layer the banks regard as shared infrastructure will be answered — and the answer, on the documented pattern, is one of three things: build it internally, pull it into an existing consortium, or start a new one. *[INFERENCE]*

The gatekeeper problem



The structural difficulty this creates for an independent vendor has four components, and they compound.

One: the customers own the competitor. The bank-owned utilities are owned and funded by the very banks any independent loan-technology vendor must sell to. An outside platform is asking its prospects to fund a rival they co-own. *[FACT AS TO OWNERSHIP; INFERENCE AS TO THE CONSEQUENCE]*

Two: they control the input the vendor needs. The banks are the administrative agents holding the golden-source loan register and borrower data. The consortia are explicitly designed to pool that data and that order flow. An independent vendor cannot assemble a competing golden source without the agents' cooperation — which they have chosen to give to their own utility. *[INFERENCE]*

Three: distribution and trust favour incumbents. Consortium strategies leverage the banks' institutional trust and distribution and are hard for outsiders to replicate. Network effects in a data or liquidity utility compound this: the platform the largest agents contribute to becomes the default. *[INFERENCE]*

Four: they can build it themselves. These institutions have the engineering capability and the precedent — internal loan-matching and bidding tools that later became consortium ventures. *[FACT]* The build-then-mutualise path captures the value while neutralising the vendor. *[INFERENCE]*

51.1 The anchor as kingmaker and gatekeeper

The consortium ventures are themselves proof of the anchor-customer playbook: they were deliberately seeded with top banks as founding customers and investors to manufacture credibility and network effects before selling to the rest of the market. *[INFERENCE]*

The logic is documented in the data platform's own launch framing: the company positioned itself as "backed by some of the biggest players and built by industry veterans," with one founder bank's chief administrative officer framing it as leading transformation of the industry's syndicated loan market. *[FACT]* The structural mechanism is that the founding banks *are* the agent banks holding the loan data, so their participation both supplies the data and signals to lenders that the platform is the emerging standard. *[INFERENCE]*

Academic work supports the general principle: incumbent-bank trust functions as an entry barrier in credit markets, which a financial-technology entrant overcomes only by borrowing an incumbent's credibility. *[FACT]* Peer-referenced proof points from a comparable institution are the highest-value signal to the next enterprise buyer. *[VENDOR CLAIM]* Business-to-business financial-technology sales cycles span six to twelve months involving multiple decision-makers — and for a bulge-bracket anchor, real timelines run longer once security, vendor-risk and legal review are added. *[VENDOR CLAIM] [INFERENCE]*

So a single anchor is a kingmaker for three compounding reasons: **data gravity** (the top banks are the agents, so they are the supply side, not merely an endorsement); **halo** (a fragmented tail of lenders and mid-tier

banks treats a top-two name as a de-risking signal); and **standard-setting** (once one or two top agents commit, the format they adopt becomes the de facto standard and holdouts face isolation cost). *[INFERENCE]*

And the same properties make the anchor dangerous. The first anchor is brutally slow to land; nobody wants to be first; procurement is heavy; the anchor knows its signature is the asset and will price and structure accordingly, potentially demanding equity, exclusivity or design-partner control; concentration risk cuts commercially, because losing the anchor loses the halo; and the anchor co-owns the competitors. *[INFERENCE]*

Kingmaker and gatekeeper are the same institution. *[INFERENCE]*

Symbiont: winning the lawsuit and dying anyway

The single most instructive failure in this market's history deserves its own chapter, because its lesson is not the one usually drawn.

The facts are in Chapter 45.1. A distributed-ledger company built a credible enterprise platform, formed a joint venture with a capital-markets data firm to move the secondary syndicated-loan market onto a shared ledger, and was destroyed when the joint-venture partner was acquired by a larger group and support was withdrawn. It sued, and won: a Delaware court held that the acquisition breached a non-competition agreement, with a headline award reported around \$70 million for its half share and a \$53 million settlement paid at the end of 2021. It filed for Chapter 11 on 1 December 2022, having exhausted its cash on legal costs, and its assets were sold for roughly \$2.6 million. *[FACT]*

The usual lesson drawn is "blockchain failed in loans." That is true and it is not the lesson.

The lesson is that in this market, a venture's fate can be determined by a transaction it is not party to. *[INFERENCE]* The company did nothing wrong technically or commercially. Its partner was acquired. The acquirer had no interest in the venture. Support evaporated, and no amount of legal vindication restored the years and the capital.


Three generalisations follow, and each has direct application to any entry strategy.

Dependency on a single partner in a consolidating market is an existential risk, not a commercial one. *[INFERENCE]* This market consolidates continuously: the data infrastructure roll-up of Chapter 12, five acquisitions of extraction companies since 2021, a \$3 billion sale process for a buy-side platform. A partner today is an acquisition target tomorrow.

Winning a dispute is not a business outcome. Litigation converts a strategic problem into a cash-burn problem with a delayed and uncertain payoff. A venture-scale balance sheet cannot absorb it. *[INFERENCE]*

The market's structure supplies no protection to an entrant that depends on it. *[INFERENCE]* There is no regulator ensuring fair access to loan data, no mandated interoperability, and no antitrust remedy at the transactional layer — the 2022 merger remedies addressed indices and commentary, not settlement or administration (Chapter 12.1). An entrant's access is contractual and revocable.

Who pays for transparency, and who loses from it



Opacity is a fork in the pipe, not a leak.

The final and deepest reason collaboration fails in this market is that transparency is not a public good here. It is a transfer.

53.1 The evidence

Academic work on the introduction of post-trade transparency in corporate bonds found "substantial reductions in the bid-ask spreads" that investors pay. *[FACT]* The counter-argument in the same literature is that transparency makes dealers reluctant to carry inventory. *[FACT]* Both can be true: transparency reduces the dealer's information rent and may reduce the dealer's willingness to provide liquidity.

In loans there is no equivalent tape. Loans are not securities and are not reported to the corporate-bond reporting facility. *[FACT]* Marks come from vendor services derived from dealer quotes. *[FACT]* Dealers therefore hold the information advantage on where paper is really trading, and the buy side pays wider, less observable spreads. *[INFERENCE]*

The trading venue that has electrified part of the market uses request-for-quote and list protocols that keep the dealer in the middle rather than anonymous all-to-all matching. *[FACT]* That is a design choice, and it is the design choice a dealer-owned venture would make. *[INFERENCE]*

53.2 The conclusion

The loan market's opacity is not a technological deficit. It is an equilibrium in which the parties with the power to end it are the parties who would bear the cost of ending it. *[INFERENCE]*

One qualification is owed to the evidence. The bid-ask finding is causal *for corporate bonds*, established around the introduction of post-trade reporting in that market. *[FACT]* Its transfer to loans is an inference, and a reasonable but not automatic one: loans are less fungible, trade far less frequently, and have a smaller and more concentrated dealer population, all of which could change the magnitude of the effect and conceivably its sign. *[INFERENCE]*

This is the most uncomfortable conclusion in the paper and it should be stated with appropriate care. It does not imply bad faith. Dealers who oppose a public tape have a legitimate argument — that transparency reduces their willingness to commit capital and therefore reduces liquidity — and that argument has support in the academic literature. *[FACT]* The point is not that anyone is behaving improperly. The point is that **on the evidence assembled here, no configuration of voluntary market action produces transparency, because the actors who would have to produce it bear its cost.** *[INFERENCE]* Chapter 87 asks what follows from that, and transparency is the clearest of the very few places in this paper where the answer is regulation.

The corollary for any entrant is direct. A thesis that depends on the market becoming more transparent is a thesis that depends on either regulatory compulsion or a change in dealer economics. Neither is in prospect: the market's own 2026 settlement-reform programme explicitly carries **no transparency mandate.** *[FACT]*

A thesis that instead offers *one institution* better visibility into *its own* positions faces none of this, because it takes nothing from anyone.

[INFERENCE] That asymmetry — between market transparency, which is blocked, and institutional visibility, which is not — is the most useful single distinction an entrant can carry out of this part.

VIII

PART EIGHT

What private credit changes

Private credit is the most discussed development in institutional lending in a decade and the most loosely described. It is variously an asset class, a business model, a regulatory arbitrage, a systemic risk and a technology opportunity, and the discussion rarely distinguishes them.

This part asks a narrow question. Holding aside whether private credit is a good place to lend money, **what does it change about the operating problem this paper has described?**

The answer has four parts. It changes the number of participants who must maintain records. It removes the agent bank from a large share of transactions without removing the agency function. It replaces mark-to-market with manager judgment, which converts a pricing question into a governance question. And it has attracted a supervisory apparatus that is now demanding, in writing, the data the market cannot produce.

It also, in 2025 and 2026, ran its first stress test. That episode is the most important recent evidence in this paper and Chapter 59 examines it closely.

Size, growth, and the arithmetic of the numbers

Precision first, because this is where most analysis of private credit goes wrong.

54.1 The range

MEASURE	FIGURE	SOURCE CLASS
2010 base	\$158 billion globally	[FACT — IMF]
Mid-2024	"nearly \$2 trillion" globally	[FACT — IMF]
Q2 2024	~\$1.34 trillion US / ~\$2 trillion globally; ~5× growth since 2009	[FACT — FEDERAL RESERVE, THE MOST CONSERVATIVE MEASURE]
End-2024	"\$1.5–2.0 trillion," heavily concentrated in a few jurisdictions	[FACT — FINANCIAL STABILITY BOARD]
2024	Surpassed \$2 trillion AUM; ~\$2.28 trillion projected 2025	[VENDOR CLAIM — PREQIN]
Start-2025	~\$3 trillion	[VENDOR CLAIM — MORGAN STANLEY]
2028	~\$2.8 trillion (Preqin) to ~\$3 trillion (Moody's)	[VENDOR CLAIM]
2029	~\$5 trillion, broad definition including asset-based	[VENDOR CLAIM — MORGAN STANLEY]

No independent source prints "\$2.5 trillion." [FACT] It is a defensible rounding of a range and an indefensible citation. The honest formulation is "roughly \$2–3 trillion depending on source and definition, with regulator estimates at the lower end."

A further caution: the higher figure sometimes quoted for combined bank and non-bank exposure is a *different and broader* metric from a different institution and must not be conflated with private-credit assets under management. [FACT]

54.2 The public window

Business development companies provide the only regularly disclosed view. Roughly \$450 billion of fair value and \$475–500 billion of assets under management as of the first quarter of 2025, up from roughly \$30 billion in 2010, with the top ten representing fifty-three per cent of the market and perpetual non-traded vehicles the fastest-growing form. *[FACT]* BDC leverage rose from roughly forty per cent to fifty-three per cent debt-to-assets. *[FACT]*

54.3 The players

Platform-level alternative assets under management: Blackstone approximately \$1.3 trillion, Apollo approximately \$938 billion, KKR approximately \$744 billion, Ares approximately \$623 billion, Brookfield approximately \$603 billion fee-bearing, with Ares' credit group at roughly \$407 billion. *[VENDOR CLAIM — CROSS-MANAGER AUM IS NOT DEFINED CONSISTENTLY AND RANKINGS SHIFT BY DEFINITION]* Blue Owl reported \$251.1 billion of total AUM at end-2024 and more than \$284 billion by 30 June 2025. *[FACT — SECURITIES FILINGS]* Blackstone's non-traded BDC passed \$64 billion in September 2024 and more than \$70 billion by year-end, described as the largest private debt fund in the world. *[VENDOR CLAIM]* Apollo reported approximately \$751 billion of total AUM at end-2024. *[FACT]*

Two consolidation events define the current structure. **BlackRock acquired HPS Investment Partners for approximately \$12 billion**, announced December 2024 and closed 1 July 2025; HPS had grown from \$34 billion in 2016 to \$148 billion by September 2024, and by 2026 the combined credit platform was reported at roughly \$190 billion. *[FACT]* **KKR completed 100 per cent ownership of Global Atlantic on 2 January 2024**, with that insurer's AUM growing from \$72 billion in 2020 to \$158 billion. *[FACT]*

The second of those points to the structural fuel. Private credit on US life-insurer balance sheets reached **\$849 billion in 2024, roughly fourteen per cent of total assets and double the level of a decade earlier**, and insurance capital now funds approximately **forty-three per cent of credit AUM at the**

top seven alternative managers, up from thirty-two per cent in 2021.

[VENDOR CLAIM] One insurer affiliated with a major manager exceeded \$470 billion of assets by June 2026. *[FACT]*

Insurance capital is the reason private credit's growth is durable rather than cyclical. *[INFERENCE]* An annuity liability is permanent capital: it does not redeem in a drawdown, it must be matched with long-dated assets, and it creates a captive, low-churn funding base for origination. The managers that own insurers have a structurally different cost of capital from those that do not.

Structurally different, not just smaller syndicates

The most common analytical error about private credit is to treat it as a syndicated loan with fewer lenders. The differences are structural.

DIMENSION	BROADLY SYNDICATED	PRIVATE CREDIT
Origination	Bank-underwritten and syndicated	Bilateral or club, directly negotiated <i>[FACT]</i>
Liquidity	Actively traded secondary market	Hold-to-maturity; limited or no secondary; bilateral unitranche deals carry five-to-seven-year lockups and reported ten-to-fifteen per cent discounts on the rare secondary exit <i>[FACT]</i>
Documentation	Trade-body model documents as a base, bespoke economics	Wholly bespoke; "tailored documentation," "unique covenants, payment schedules and waterfall structures" <i>[FACT/VENDOR CLAIM]</i>
Covenants	Over ninety per cent of BSL is covenant-lite <i>[FACT]</i>	Historically maintenance covenants — but eroding <i>[FACT]</i>
Valuation	Quote-derived marks from vendor services	Manager-determined fair value, not mark-to-market <i>[FACT]</i>
Agent	Bank agent with a servicing platform	Often no bank agent; non-bank agents and third-party administrators <i>[FACT]</i>
Lender base	Large, uncoordinated, may not be able to coordinate	Small, aligned, amendment consents faster <i>[FACT]</i>

Three of these deserve elaboration.

The covenant advantage is eroding. A rating agency finds that "flexible structuring continues to reign in private credit," meaning the documentation advantage over the syndicated market is narrowing as competition intensifies. *[FACT]* The differentiator that justified private credit's underwriting premium is being competed away.

The valuation difference is the load-bearing one. In the syndicated market, a position's value is a number someone else produces from dealer quotes; the manager can disagree with it but cannot make it go away. In private credit the manager produces the number. That is not a criticism — there is

no market price to use — but it converts a pricing question into a **governance** question: what process, what inputs, what evidence, what review. *[INFERENCE]* Chapter 56 shows what supervisors have concluded about that.

The absence of a bank agent does not eliminate the agency function.

Someone must still maintain the register, calculate interest, issue notices, process transfers and distribute cash. In private credit that someone is the manager, a non-bank agent, or a third-party administrator — often without the servicing infrastructure a bank agent has. *[INFERENCE]* Private credit removed the agent bank; it did not remove the work.

Valuation, opacity, and the arrival of the supervisor

Between 2024 and 2026 the official sector converged on private credit with unusual consistency, and the specific language matters more than the fact of attention.

The International Monetary Fund, April 2024: borrowers are "smaller and riskier"; the sector "has never experienced a severe downturn at its current size"; vulnerabilities arise from "fragile borrowers, a growing share of semi-liquid vehicles, multiple layers of leverage, stale and potentially subjective valuations, and unclear connections between participants." The verdict: it could "become systemic" if opacity persists. *[FACT]*

The Bank of England: valuation is "inherently subjective... considerable discretion"; data gaps mean banks "cannot systematically aggregate combined credit and counterparty risk exposures." Private-credit funds have been included in its system-wide exploratory stress-test exercise. *[FACT]*

The Financial Stability Board, 6 May 2026, in a formal report on vulnerabilities in private credit: four vulnerability clusters — bank interconnections; borrower credit quality and valuation opacity; concentration and liquidity mismatch; and **data gaps**. The report sizes the market at \$1.5–2.0 trillion and bank exposure at approximately \$220 billion drawn and undrawn, against commercial estimates of \$270–500 billion — a range that itself illustrates the data problem. Recommended actions: close data gaps, harmonise definitions, deepen interlinkage analysis. *[FACT]*

The Securities and Exchange Commission's 2026 examination priorities target private-credit valuation, liquidity management, lock-ups and disclosure, with examinations to assess fair-valuation methods and controls for illiquid assets. *[FACT]* Form PF compliance was extended to 1

October 2026, with a proposed April 2026 amendment that would raise the reporting threshold from \$150 million to \$1 billion in AUM, reducing the filer population. *[FACT]*

Ratings agencies report that "hidden leverage — through PIK debt and NAV lending — is becoming more prevalent," with weaker underwriting and reduced visibility. *[VENDOR CLAIM]*

Litigation followed. Business development company lawsuits in 2026 allege misstated net asset values and inadequate valuation processes. *[FACT]*

The consistency of the language across four institutions with different mandates is the finding. **"Valuation opacity" and "data gaps" are not vendor talking points. They are the terms in which the official sector has framed its concern, in writing, in formal publications.** *[FACT]*

The strategic implication is specific and often overstated, so it should be put carefully. This converts governance and provenance from an efficiency proposition into a compliance proposition, which is a materially better sales position because compliance budgets are less discretionary than efficiency budgets. *[INFERENCE]* It does **not** mean supervisors will mandate any particular technology; none has, and none is proposing to. *[FACT]* What it means is that the buyer of a governance capability can now justify the purchase by pointing to a regulator's published concern rather than to a productivity estimate. *[INFERENCE]*

The operating layer: spreadsheets at scale

The characterisation of private credit's operating reality is remarkably consistent across sources that have no reason to agree.

From an operations software vendor: private credit is "built around bilateral lending relationships, tailored documentation, and buy-and-hold structures... each deal comes with unique covenants, payment schedules, and waterfall structures," and data arrives as "PDFs, emails and spreadsheets... each structured differently and requiring manual reading."

[VENDOR CLAIM]

From a fund-administration provider: "A spreadsheet that worked for 20 loans breaks down under the weight of 200," producing data errors, missed billing cycles and no real-time portfolio visibility. *[VENDOR CLAIM]*

From a large technology and services provider: disconnected data across investment, finance, investor-relations and compliance functions, where "normalising unstructured information — credit agreements, agent bank notices, portfolio company financials — remains the sector's biggest operational challenge." *[VENDOR CLAIM]*

From the business press: "Private Credit's Spreadsheet Problem Draws An AI Servicing Push." *[FACT — HEADLINE]*

From the Financial Stability Board, which sells nothing: regulators face "data challenges when monitoring exposures." *[FACT]*

When four vendors, the trade press and a global supervisor describe the same gap, the gap exists. *[INFERENCE]* The vendor characterisations should still be read as vendor characterisations — they are selling the fix — but the supervisory corroboration removes the possibility that the problem is manufactured.

What specifically breaks, and why the threshold is real. The claim that a spreadsheet fails somewhere between twenty and two hundred loans is a vendor's, but the mechanism behind it is structural and can be stated without relying on the vendor. *[INFERENCE]* Five things scale badly at once.

The unique terms. Each deal carries its own covenant definitions, payment schedule, PIK election mechanics and waterfall. In a syndicated facility much of this is at least drafted from recognisable market forms; in direct lending it is negotiated bilaterally, which means the number of distinct interpretive objects grows linearly with the portfolio rather than being amortised across it.

The calculation surface. Interest accruals, fee accruals, PIK capitalisation and amortisation each have to be computed per deal per period. A spreadsheet does this correctly until someone changes a formula, and there is no control that detects it.

The reporting cadence. Borrower financials and compliance certificates arrive as documents, on staggered dates, in the format the borrower chooses. Someone reads each one and re-keys it. That workload is proportional to deals multiplied by periods, and it is the single largest recurring manual task in the sector. *[INFERENCE]*

The investor side. Capital calls, distributions, allocations, side-letter terms and bespoke investor reporting run on the same spreadsheets as the asset side, so an error on one side propagates to the other with nothing in between to catch it.

The valuation obligation. Chapter 56's supervisory attention lands here, and it lands on a process rather than an answer: what was used, by whom, on what basis, reviewed by whom, consistently with what was done last quarter. An operating layer that cannot produce that record cannot demonstrate the governance, whatever the valuations are.

None of that is exotic, and none of it requires a system to do badly. It requires a system to do *provably*, which is a different specification and is the one the sector is now being asked for. *[INFERENCE]*

The comparison that frames the opportunity: even *syndicated* loan operations are considered antiquated. *[VENDOR CLAIM]* Private credit is a rung below that in standardisation. *[INFERENCE]* It has more bespoke documentation, less market infrastructure, no trade-body settlement conventions of comparable reach, and a lender base that has grown faster than its operating capability.

The capital dynamic

Three reinforcing money flows shape the current structure, and the third is the most interesting.

Retail and institutional fundraising. Through the third quarter of 2025, non-traded BDCs raised \$167 billion and interval funds \$126 billion, the latter opening the asset class beyond accredited investors. *[VENDOR CLAIM]*

Insurance balance-sheet capital, covered in Chapter 54.3.

Banks funding private credit. This is the structural novelty. Federal Reserve research shows bank committed credit lines to private-credit vehicles grew from roughly \$8 billion in 2013 to approximately **\$95 billion by the fourth quarter of 2024**, with about \$56 billion utilised: \$79 billion in revolvers — \$49 billion to BDCs and \$30 billion to private debt funds — plus \$16 billion of term loans, with the largest US banks holding roughly sixty per cent. *[FACT]*

Beyond lending to the sector, banks have entered joint ventures with it. Citi and Apollo announced a **\$25 billion** direct-lending programme in September 2024 with an insurer and a sovereign investor participating and a first-year origination target of \$5 billion; Citi and BlackRock's HPS agreed in May 2026 to collaborate on European direct lending, up to **€15 billion** over five years. JPMorgan increased a direct-lending commitment to **\$50 billion** in February 2025, plus approximately \$15 billion from co-lenders, and invited direct-lending partners. Wells Fargo and a private-equity partner launched a \$5 billion fund. *[FACT]*

Managers have also begun trying to manufacture secondary liquidity where none exists structurally, with a major manager teaming with two large banks to trade private credit in May 2025. *[FACT]* That is a tell: the illiquidity is being actively worked because it has become a constraint. *[INFERENCE]*

The structural reading: the bank/non-bank boundary that the post-crisis regulation created is being deliberately dissolved by both sides. *[INFERENCE]* Banks originate and distribute; managers hold; insurance funds the holding; and the loan sits in a structure that is neither purely bank nor purely fund.

For the operating problem, this matters more than the credit question. A **hybrid origination structure** means a single facility may touch a bank's servicing platform, a manager's portfolio system, an insurer's statutory accounting, a fund administrator's books and a rating agency's model — with no single party holding an authoritative record across all of them.

[INFERENCE] The duplication census of Chapter 25 gets longer, not shorter.

The stress test: First Brands, Tricolor, gating, and PIK

In late 2025 and through 2026 private credit experienced its first meaningful stress. What broke, and how, is the most important recent evidence in this paper.

59.1 The collateral frauds

First Brands filed in **September 2025** with **liabilities** its petition put in a range of **\$10 billion to \$50 billion**; the Justice Department's indictment cites more than \$9 billion, and analysts' totals including off-balance-sheet financing reach roughly \$11.6 billion. **Tricolor** collapsed with **\$900 million-plus**. Both were driven by **double-pledged collateral and fabricated receivables**. Founders were criminally charged — Tricolor in December 2025, First Brands in January 2026. *[FACT]*

A note on what these numbers are, because they are routinely misused — including in the first draft of this paper. Five different quantities circulate around these two events and they are not interchangeable:

QUANTITY	WHAT IT MEANS	DOCUMENTED FIGURES
Liabilities	What the failed company owed, in total, to all creditors	First Brands: \$10–50bn petition range; over \$9bn per the DOJ indictment <i>[FACT]</i>
Collapse size	Loose shorthand for the scale of the failure	Tricolor \$900m+ <i>[FACT]</i>
Gross exposure	What a given lender had at risk before recovery	Jefferies ~\$715m via an affiliated vehicle; a UBS unit >\$500m <i>[FACT]</i>
Charge-off	What a lender wrote off against earnings	JPMorgan \$170m on Tricolor <i>[FACT]</i>
Realized loss	What a lender ultimately lost after recovery	Jefferies ~\$10m hit after a prior \$30m, having written its First Brands exposure to zero <i>[FACT]</i>

Adding the first two together and calling the sum a loss is wrong, and it is wrong by a very large factor: liabilities include everything the estate owes and say nothing about lender recovery, while the publicly documented realized losses at named institutions are two to three orders of magnitude smaller than the headline. This paper therefore does not use an aggregate loss figure for these events, and Appendix D records the error. What the events establish is the *failure mode* and its reach across named institutions — which is what the argument actually needs. Jefferies was sued by investors and came under regulatory scrutiny. *[FACT]*

The same discipline governs a second class of figure that circulates around these events: the difference between what is **alleged** in a complaint, what is **charged** in an indictment, and what has been **found** by a court. Tricolor's collateral shortfall — approximately \$2.2 billion pledged against approximately \$1.4 billion of actual collateral, and roughly 29,000 of some 70,000 active loans sharing identical vehicle identification numbers — rests on allegations and reporting, not adjudication: the \$2.2 billion figure is a prosecutors' estimate in the criminal case, and the 29,000 figure was reported by Bloomberg from the bankruptcy record. *[FACT — THAT THESE ARE THE ALLEGATIONS]* First Brands' fabricated or inflated receivables, reported at between roughly \$2.3 billion and \$2.7 billion depending on which measure is cited, come from estate professionals and reporting rather than from a judgment. *[FACT — AS REPORTED]* The criminal proceedings are at the charging stage in both cases. Appendix F sets out every figure used in this paper for these events with its class marked, and adds no figure to any other.

A further reported UK property-lender failure with an alleged collateral shortfall in excess of \$1 billion surfaced in one outlet and **could not be independently corroborated**; it is noted and not relied upon. *[UNKNOWN]*

59.2 What kind of failure this was

This is the analytically important part, and it is frequently misread.

These were **not credit-underwriting failures** in the ordinary sense — a lender lending to a borrower that subsequently deteriorated. They were **collateral-tracking and data-integrity failures**: the same asset financed twice, and receivables that did not exist. *[INFERENCE]* The lenders were not wrong about the borrower's prospects. They were wrong about what they had a claim on.

That distinction determines what, if anything, technology could have addressed. A better credit model would not have helped, because the lenders were not wrong about the borrower's prospects. Whether a system capable of establishing that a given receivable existed and had been pledged to exactly one lender **would** have prevented these particular frauds is a counterfactual about events that did not happen, and this paper states it as such: it is plausible on the mechanism, and it is not demonstrated.

[HYPOTHESIS] The weaker claim is safe and sufficient: **these were failures of a kind that no party in the structure was authoritative about** (Chapter 69), which is a different and more addressable defect than a failure of credit judgment. *[INFERENCE]*

The consensus framing in the market has been "idiosyncratic fraud, not systemic," which tempers the crisis narrative. *[FACT]* It also sharpens the operational one: idiosyncratic frauds succeed when the operational layer cannot see the truth of the collateral, and the operational layer in private credit is, on the evidence of Chapter 57, spreadsheets and PDFs. *[INFERENCE]*

59.3 The liquidity stress

Separately and concurrently, the retail channel seized.

A major manager gated a retail debt fund in February 2026 and capped redemptions at five per cent in April 2026; one of its BDCs switched from quarterly tenders to return-of-capital distributions; it sold \$1.4 billion of loans at 99.7 per cent of par. Another large manager also limited redemptions. A flagship corporate lending fund imposed a five per cent withdrawal limit in March 2026. *[FACT]*

Fourth-quarter 2025 redemptions at BDCs above \$1 billion reached \$2.9 billion, up roughly 200 per cent quarter on quarter, with average non-traded BDC redemptions rising to 4.8 per cent of net asset value from 1.6 per cent. *[VENDOR CLAIM]*

Credit metrics deteriorated. A broad measure put the US private-credit default rate at a record **6.0 per cent** in April 2026, with roughly sixty-five per cent of 2025 defaults attributed to distressed restructurings. BDC payment-in-kind share rose from approximately six per cent in 2022 to roughly **ten per cent** by early 2026, and "bad PIK" — deferral due to liquidity strain rather than by design — reached 6.4 per cent of loans, roughly three times the 2021 level. *[FACT/VENDOR CLAIM]*

The selling of loans at 99.7 per cent of par while marks were held at par is precisely the valuation tension the supervisors named. *[INFERENCE]*

59.4 What the episode establishes

It strengthens the operational case and complicates the growth case. *[INFERENCE]*

The operational case is strengthened because the failures were of a kind an operating and governance layer plausibly addresses: collateral integrity, position verification, and the ability to establish whether a claimed asset is what it is claimed to be. This is the strongest evidence in the paper that the market's data layer carries real, recent, attributable costs at named institutions — with the arithmetic caveat of \$59.1 firmly attached, since the honest figures are exposures and charge-offs rather than an aggregated loss.

The growth case is complicated because a gating retail channel and rising defaults mean inflows may slow. The wedge argument therefore should not rest on "private credit is growing fast." It should rest on **the governance requirement, which grows whether or not inflows do — and grows faster in a downturn.** *[INFERENCE]*

There is a further caution. A downturn is a double-edged catalyst: it exposes the operational and valuation weakness, creating urgency, and it freezes budgets. *[INFERENCE]* Which effect dominates in any given institution is not predictable from the outside. *[UNKNOWN]*

Does private credit repeat the syndicated market's infrastructure history?

The strategically decisive question about private credit is whether it is an early-stage version of the syndicated market — in which case its infrastructure will consolidate along the same path — or a structurally different market that will develop differently.

The paper's honest answer is that both arguments are strong and the evidence currently favours convergence.

60.1 The case that it repeats

The players are the same, or are becoming so. A bank consortium data platform added a mega-manager to its cap table in April 2026 and is explicitly courting the buy side and private credit. *[FACT]* The dominant data provider extended its identifier franchise to private-credit assets in June 2025. *[FACT]* The dominant servicing vendor is aiming its platform at private credit with a dedicated positioning. *[FACT]* The buy-side platform serving CLO managers is being sold at up to \$3 billion. *[FACT]* Private-credit CLOs already represent roughly nineteen per cent of CLO issuance. *[FACT]*

The economics are the same. Insurance and permanent capital want scale; scale wants standardisation; standardisation wants shared infrastructure; shared infrastructure in this market is built by consortia (Part Seven). *[INFERENCE]*

The lending is converging. Bank–manager joint ventures blur the origination boundary; repricings flow both ways between the syndicated and private markets; and a rating agency notes the private credit boom narrowing the gap to the syndicated market. *[FACT]*

60.2 The case that it does not

There is no agent bank to organise around. The syndicated market's infrastructure is built on the agent as a data source. A market where the manager is the agent has no analogous chokepoint to pool. *[INFERENCE]*

There is no secondary market to standardise for. The syndicated market's trade bodies were founded to grow secondary liquidity, and its documentation, identifiers and settlement conventions all serve trading. A hold-to-maturity market has weaker incentives to standardise. *[INFERENCE]*

The buyer population is fragmented. The syndicated market has a top-four that clears roughly thirty per cent of volume. Private credit has a few mega-managers and a long tail of boutiques, which is a harder go-to-market and a weaker basis for standard-setting. *[INFERENCE]*

Confidentiality is a feature. Borrowers choose private credit partly for confidentiality. *[FACT]* Infrastructure that increases transparency runs against a stated customer preference — which is a different obstacle from the syndicated market's, where transparency is resisted by intermediaries rather than by borrowers. *[INFERENCE]*

60.3 The judgment

On the current evidence, convergence is more likely than divergence, and the convergence is being driven deliberately by incumbents rather than emerging organically. *[INFERENCE]*

The decisive datapoint is the April 2026 consortium round with a mega-manager investor and an explicit private-credit and buy-side push. *[FACT]* That is the syndicated market's institutional form — the bank-owned utility — being extended into private credit by the same actors, with a manager now inside the tent.

The window in which private credit's operating layer is genuinely unclaimed is therefore narrower than the growth statistics suggest, and it is narrowing by deliberate action rather than by drift. *[INFERENCE]* Chapter 78

treats this as the central timing constraint on any entry thesis.

Private credit therefore intensifies several problems already present in syndicated lending while changing where they sit.

Bespoke interpretation remains. Operational state remains distributed. Valuation becomes more judgment-dependent. Collateral verification becomes more consequential. The agency function persists even where the traditional agent bank does not. And supervisors increasingly demand evidence about processes whose underlying information is difficult to assemble. *[INFERENCE]*

None of those observations establishes that AI solves the problem.

They do establish why a change in the economics of interpretation would matter.

IX

PART NINE

What ai changes, and what it does not

The central question of this paper is what has changed between the last generation of modernization attempts and 2026. The candidate answer that everyone reaches for is artificial intelligence, and the candidate answer is usually asserted rather than argued.

This part argues it, in both directions. It establishes what was not possible in 2022 and is possible now, with citations to the research rather than to marketing. It then spends more space on the limits than on the capabilities, because the limits are where the strategic content is. A reader who takes only one thing from this part should take Chapter 63, on the evaluation problem, and Chapter 65, on what must remain deterministic.

The conclusion, stated in advance: **AI changes exactly one of the five complexity categories from Chapter 26, and the market's historical failures were in two different ones.** That is a real change and it is not obviously a sufficient one.

The 2022 baseline

To see what changed, it is necessary to be honest about what was and was not possible before.

In 2022, a firm wanting to convert a 200-page bespoke credit agreement into structured servicing data had three options.

Have a person read it. Accurate, expensive, slow, and non-scalable. This is what the market did for the deals that mattered.

Configure a rules engine. A human reads the agreement and encodes its terms as structured configuration in a servicing platform. Accurate if done correctly, expensive, slow, and — critically — **brittle**, because the encoding is per-deal and must be redone when the agreement is amended. The documented consequence is that operations shops fall back to spreadsheets when the encoding burden exceeds the platform's value. *[FACT THAT THE FALLBACK OCCURS; INFERENCE AS TO CAUSE]*

Use optical character recognition and template extraction. This worked for standardised documents with predictable layouts. Credit agreements are neither. Template-based extraction requires the target document to resemble the template, and the entire point of Chapter 36 is that the commercially significant provisions are drafted uniquely. *[INFERENCE]*

So the market's position in 2022 was: reading the agreements was possible but not economical, and no technology existed that made it economical.

The interpretation category — Category 2 in Chapter 26 — was written off. *[INFERENCE]*

It is worth registering how much of the market's operating design adapted to that constraint. Institutions relied on notices rather than continuously rereading contracts; relied on the agent's interpretation rather than independently reproducing it; maintained spreadsheets where bespoke

terms exceeded the practical encoding capacity of core systems; and trained operations specialists over long periods to understand documents whose consequences were difficult to represent systematically. *[INFERENCE]*

The architecture was not irrational. It was an adaptation to the economics of interpretation.

If those economics change, some of the architecture can be reconsidered without assuming that the contract, authority structure or servicing core changes with it. *[INFERENCE]*

What is genuinely new

Between 2023 and 2026 a body of peer-reviewed and industrial evidence accumulated on exactly the capability the market wrote off. It is worth citing precisely rather than gesturing at it.

Auto-documentation of undocumented legacy systems. A peer-reviewed study presented at a software-engineering workshop at a top-tier 2025 conference found that large-language-model-generated comments for legacy mainframe and MUMPS code were "generally hallucination-free, complete, readable, and useful compared to ground-truth comments." *[FACT]*

Specification and behaviour extraction as a named research direction. Published frameworks explicitly aim at "converting legacy software into operational specifications" and at "preserving business logic with behavioural specification graphs." *[FACT]*

Production tooling. A major vendor's mainframe code assistant, generally available from 2023 and agentic in later versions, generates "natural language explanations of COBOL code to identify the right modernization strategy," and provides a data dictionary that embeds "client-specific meaning directly into the AI workflow." *[VENDOR CLAIM]*

Independent institutional validation. A federally funded research organisation published guidance on legacy IT modernization with AI in June 2025. *[FACT]*

Application to the market's own documents. Practitioners cite vision-based large language models being used to parse "semi-structured and unstructured documents and PDFs" for loan monitoring. *[FACT –*

PRACTITIONER REPORTING]

Application to expert knowledge capture. Academic work now frames generative AI as the vehicle to externalise apprenticeship knowledge, with a 2026 paper on making tacit knowledge explicit in apprenticeship systems

and a retrieval-augmented architecture for expert knowledge preservation.

[FACT] Practitioner guidance describes AI transcribing and structuring expert interviews to capture "the nuance of decision-making, not just final results." [FACT]

62.1 What this actually establishes

Precision matters here, because the gap between what the research establishes and what is claimed for it is wide.

Established: the cost of producing a plausible, readable, largely accurate description of an undocumented system or an unstructured document has fallen by a large factor. [FACT]

Established: this applies to the specific problem of a legacy system whose behaviour no living person fully understands — which is the Chapter 66 problem, and which is the incumbent servicing platform's situation as its implementation expertise disperses (Chapter 34). [INFERENCE]

Not established: that the output is correct at the grain required for a financial calculation. [FACT — SEE CHAPTER 63]

Not established: that any of this addresses coordination or institutional control. [INFERENCE]


The honest formulation is that **AI has made Category 2 economical for the first time.** [INFERENCE] Reading every credit agreement in a portfolio, and reading every notice, is now a cost that a mid-sized institution could contemplate. That is a genuine change and it is the only genuine change. Everything else in the enthusiasm is either an extrapolation from it or a restatement of capabilities that existed before.

Put less taxonomically: **AI changed the economics of interpretation.**

It changed what an institution can economically know about documents it already possesses. It did not thereby change what that institution is contractually entitled to decide. [INFERENCE]

FIGURE 8

What changed, and what did not



Interpretation is the one category of this market's complexity whose economics moved. Calculation improved incrementally; authority, consent and external evidence did not move at all. The adoption metrics in the left column are a vendor's account of its own workflows and carry that class in Chapter 47. [VENDOR CLAIM]

The evaluation problem

This chapter contains the most important limitation in the paper, and it is one that promotional material about AI in credit uniformly omits.

The same peer-reviewed study that found generated comments "hallucination-free, complete, readable and useful" also found that **"no automated metrics strongly correlate with comment quality"** and that **"human evaluation remains necessary."** *[FACT]*

Read those two findings together. The technology can produce good output. **We cannot cheaply tell whether it did.**

A useful distinction follows:

Hallucination-free is weaker than financially correct.

A system can invent nothing and still omit a margin tier, apply the correct provision to the wrong tranche, miss an interacting definition, use the wrong agreement version, or otherwise produce an interpretation that is textually defensible and financially wrong.

In this domain, plausibility is not the acceptance criterion.

For a summarisation task this is an inconvenience. For a servicing task it is close to disqualifying on its own, and the reason is arithmetic rather than philosophical.

Suppose an extraction system reads margin grids from credit agreements with ninety-nine per cent accuracy — far better than any published result on documents of this complexity. A portfolio of a thousand facilities then contains ten wrong margin grids. Each wrong grid produces a wrong interest calculation on every payment date for the life of the facility. Each wrong calculation produces a wrong payment to a set of lenders, a break in every downstream system, and — as Revlon established — a legal position in which recovering the money is not guaranteed. *[INFERENCE]*

The problem is not the error rate. The problem is that **the errors are silent**. A wrong margin grid does not announce itself. It produces a perfectly plausible number that is wrong, in a system where the counterparty's independent check is against data derived from the same source. *[INFERENCE]*

63.1 Why "hallucination-free" is the wrong bar

The research finding is that generated documentation is hallucination-free — meaning it does not invent things that are not in the source. That is a meaningful and hard-won property.

It is a far weaker bar than **financial correctness**. A description of a margin grid can contain nothing invented and still be wrong: it can omit a tier, misread which ratio the grid keys off, apply the grid to the wrong tranche, or miss an interaction with a separate defined term thirty pages away. None of those is a hallucination. All of them are wrong answers. *[INFERENCE]*

63.2 The consequence

AI can propose. Something else must prove. *[INFERENCE]*

That is not a rhetorical formulation; it is an architectural constraint with a name. The engineering pattern that supplies the proof is the **parallel run**: call both the old and the new implementation, compare the results, and treat only one as the source of truth at any time, verifying that the new implementation gives the same answers within acceptable bounds. *[FACT —*

THIS IS A DOCUMENTED, NAMED PATTERN WITH PUBLISHED AUTHORSHIP]

The relationship between the two is precise and worth stating carefully: **AI proposes the model; the parallel run proves it.** *[INFERENCE]* Neither is sufficient alone. The parallel run without AI is the expensive manual encoding of 2022. AI without the parallel run is an unverified financial calculation, which is not a product anyone in this market can buy.

And the parallel run has its own precondition that is easy to miss. **It presupposes you can compute the same answer the incumbent computes**, which presupposes you have already correctly reverse-engineered the

incumbent's business logic — the fee accruals, the syndicate allocation, the waterfall. *[FACT]* The parallel run validates the model; it does not create it. So the two techniques are mutually dependent, and the loop closes only if the AI's proposal is close enough to be worth verifying. *[INFERENCE]*

Why "AI reads contracts" is not a business

The market has already reached this conclusion, and the evidence is in the transaction record.

Chapter 47 lists the extraction field: deep-research platforms, document processors, clause extractors, alternative-asset ingestion tools, origination systems, securitisation document readers, valuation-first extraction platforms, and general legal and research AI. A dozen or more entrants.

The consolidation record since 2021: Finley acquired by Concord in February 2026; Solvas acquired by Alter Domus in May 2023; Lumonic acquired by Morningstar in March 2025; Kira Systems inside Litera since 2021; Evisort acquired by Workday. *[FACT]*

When incumbents buy a capability rather than compete with it, the capability is being absorbed into the baseline. *[INFERENCE]*

The most telling instance is Lumonic, acquired by Morningstar and positioned as "the only AI-native private-credit spreading, end-to-end, **with source-cell traceability.**" *[FACT]* Source-cell traceability is provenance. It is the differentiator any governance-oriented thesis would claim, and a data incumbent with distribution already owns it.

Three conclusions.

Extraction is table stakes. *[INFERENCE]* A product whose differentiation is that it reads credit agreements with AI is differentiated on a capability that a dozen firms have and several incumbents have bought.


Provenance alone is not enough either. *[INFERENCE]* It is better than raw extraction, and it is also already claimed.

The remaining differentiators are the ones that are hard for structural rather than technical reasons. *[INFERENCE]* Being connected to the authoritative source. Being able to reconcile continuously rather than

extract once. Being able to act on the result rather than merely report it. Having the contractual right to sit where the data is. Those are Category 3 and Category 5 problems (Chapter 26), and no amount of model capability supplies them.

FIGURE 9

Analysis is not authority



A system can identify the voting rule, the affected lenders and the threshold, and can draft the documentation. It cannot cast a lender's vote, waive a right, or override a sacred right — those belong to the parties named in the contract. Consent thresholds shown reflect common market practice rather than any single agreement. *[FACT]*

What must remain deterministic

DETERMINISTIC, ALWAYS

- Monetary calculations
- Balances and positions
- The register
- Allocations
- Ledger entries
- Entitlements
- Transaction state

PROBABILISTIC, WITH VERIFICATION

- Reading a bespoke agreement
- Classifying a notice
- Resolving semantic ambiguity
- Explaining a discrepancy
- Drafting
- Triaging exceptions
- Capturing expert knowledge

Any credible architecture in this market draws a hard line, and the line is worth drawing explicitly because it is where a great deal of enthusiasm should stop.

65.1 The deterministic side

The following must never be produced by a probabilistic system:

- **Monetary calculations:** interest accruals, fee computations, compounding, day-count conventions.
- **Balances and positions:** principal outstanding, commitment, undrawn amount.
- **The register:** who owns what, as of when.
- **Allocations:** the split of a payment across a syndicate.
- **Ledger entries and accounting state.**
- **Entitlements:** who is entitled to vote, to receive, to consent.
- **Transaction state:** whether a trade has settled, whether a consent has been given.

The reason is not that models are bad at arithmetic. It is that these values are **legally operative**. A payment allocation is not an estimate; it is a determination that moves money and creates or extinguishes obligations. The system must be able to say not only what the number is, but *why* — deterministically, reproducibly, and in a form that survives audit and litigation. *[INFERENCE]*

Revlon supplies the empirical support. The dispute turned on whether the lenders were entitled to keep money sent in error, and the courts resolved it by reference to whether the debt was presently due — a determination about contractual state. *[FACT]* A system whose answer to "was this payment correct?" is a confidence score has no place in that chain. *[INFERENCE]*

65.2 The probabilistic side

The following are appropriate for a model, with verification:

- **Reading a bespoke agreement** to propose a structured representation of its terms.
- **Classifying an unstructured notice** by type and extracting its fields.
- **Resolving semantic ambiguity** — determining whether this sender's "facility fee" means a commitment fee or a utilisation fee.
- **Explaining a discrepancy** between two records and proposing the likely cause.
- **Drafting** a communication, a summary, a memorandum.
- **Triaging** exceptions by likely severity and cause.
- **Capturing expert knowledge** from practitioners and making it queryable.

Note what unites these: each produces an **input to a decision** rather than the decision, and each has a human or deterministic check downstream. *[INFERENCE]*

65.3 The boundary is the product

The strategically important observation is that **the boundary between these two lists is itself the hardest engineering and the most defensible position.** *[INFERENCE]*

An extraction tool sits entirely on the probabilistic side and produces suggestions. A servicing platform sits entirely on the deterministic side and requires someone to have already done the interpretation. The valuable and difficult thing is the mechanism that carries a proposed interpretation across the boundary with enough evidence attached that a deterministic

system, an auditor and a supervisor will accept it: the clause it came from, the confidence, the reviewer who approved it, the version of the agreement, the date, and the record of every subsequent change. *[INFERENCE]*

That is a governance and provenance problem, not a model problem. It is also, notably, the thing that a data incumbent has already begun to market under the name "source-cell traceability." *[FACT]* The position is defensible in principle and contested in practice.

The important architectural object is therefore not the model alone.


It is the crossing between probabilistic interpretation and deterministic financial state.

That crossing is where provenance, evidence, review, version control and authority must be preserved. *[INFERENCE]*

AI can propose. Deterministic systems can calculate and execute. Verification establishes evidence. Governance controls the crossing.

FIGURE 10

The crossing between interpretation and execution



Probabilistic interpretation proposes; deterministic systems calculate and execute; verification establishes the evidence; governance controls the crossing between them. §65.3 argues that this crossing, not the model, is the hard engineering and the defensible position. [INFERENCE]

The workforce collision

Two trends are converging on this market's human capability, and their interaction is more consequential than either alone.

66.1 The knowledge cliff

Loan operations is learned by doing. The remedy prescribed for its training problem — shadowing, reverse-shadowing, peer training, starting with small caseloads and "advancing to full loan administration capacity as they learn the nuances of servicing different loans" — is a textbook description of apprenticeship transfer of tacit knowledge. *[VENDOR CLAIM]*

The measured figures, all from a single consultancy source and repeatedly quoted: **time to productivity for new hires of one to two years**, against **attrition of twenty to seventy per cent among onshore, US-centric loan operations talent**, "leading to a substantial loss of institutional knowledge." Syndicated loans were approximately ten per cent of worldwide corporate-banking revenue in 2022, so the function is material rather than niche. *[VENDOR CLAIM — THE PRIMARY SOURCE COULD NOT BE LOADED DIRECTLY IN RESEARCH AND THESE FIGURES ARE QUOTED FROM SEARCH EXTRACTS; THEY SHOULD BE VERIFIED BEFORE BEING RELIED UPON IN ANY EXTERNAL DOCUMENT]*

There is no public dataset reporting the age distribution or median age of loan-operations staff as a discrete occupation. *[UNKNOWN]* Official statistics publish age data for broad financial-activities categories, not this sub-function. The workforce-aging claim for loan operations rests on the training and attrition proxies above, on qualitative trade commentary, and on an analogy — and the paper states this rather than implying a hard statistic exists.

The analogy is the COBOL knowledge cliff, and it is the best-documented parallel available. The average COBOL developer is commonly cited at fifty-five to fifty-eight years old, with roughly ten per cent of the COBOL

workforce retiring each year, and the observation that "by 2030 the generation who built these systems in the 1970s–80s will be completely out of the workforce." *[FACT – WIDELY REPORTED, THOUGH THE FIGURES TRACE TO 2017*

REPORTING AND ARE REPEATEDLY RE-QUOTED SECONDHAND; THEY SHOULD BE CITED AS

INDUSTRY CONSENSUS WITH A 2017 ORIGIN RATHER THAN AS FRESH MEASUREMENT] The

scale of dependency is cited at roughly 220 billion lines of COBOL in financial systems, approximately \$3 trillion a day in transactions, roughly forty-three per cent of banking systems, and about ninety-five per cent of ATM transactions; a later industry census revised the total live COBOL figure upward to roughly 800 billion lines. *[FACT – SAME 2017-ORIGIN CAVEAT]*

Replacement failure is explicit: "COBOL experts are retiring, and their replacements don't have the education or experience," with banks rehiring retirees as contractors at \$100–300-plus an hour. *[FACT]*

And AI is already the named remediation for that analogous problem. *[FACT]*

66.2 The layoffs

The second trend arrived in 2026 and points the same direction from the opposite end.

Bank layoffs in 2026 have been concentrated in back and middle office and attributed to AI. Reported figures include roughly 74,855 finance and banking layoffs across nineteen firms as of 2 September 2026; HSBC at approximately 20,000, or 9.5 per cent; Citi at roughly 20,000 structural; Standard Chartered at 7,800 back-office and corporate roles; and JPMorgan citing AI reducing some areas by thirty to forty per cent. AI has been cited in roughly 205,000 US job losses concentrated in finance back offices. *[FACT – DIRECTIONAL; SEVERAL OF THESE COUNTS COME FROM AGGREGATOR SITES AND THE EXACT NUMBERS SHOULD BE TREATED AS INDICATIVE RATHER THAN AUDITED]*

Related survey projections: nearly eight in ten financial-services executives expect their workforce to shrink by at least twenty per cent over five years as AI absorbs processing, back-office and compliance work *[VENDOR CLAIM]*;

and Wall Street is projected to eliminate roughly 200,000 roles over three to five years as AI absorbs "entry-level processing, back-office reconciliation, and routine compliance" *[VENDOR CLAIM — A SURVEY PROJECTION, AND THE REPORTING OUTLET ITSELF CAUTIONS THAT NEAR-TERM AI LAYOFFS ARE SO FAR "MORE HYPE THAN TAKEOVER"]*.

66.3 The collision

Put the two together and the picture is specific.

The people who understand how loan operations actually works are leaving through the front door — retiring and turning over at reported rates of twenty to seventy per cent — and being cut through the back door in reductions their institutions attribute to AI, while the process they were operating remains substantially manual and the knowledge required to operate it takes a reported one to two years to acquire. *[INFERENCE, RESTING ON FIGURES THE PAPER HAS LABELLED AS VENDOR-SOURCED AND INDICATIVE RATHER THAN MEASURED]* The attribution of the layoffs to AI is the institutions' own; whether AI is the cause or the occasion is not something the public record settles. *[UNKNOWN]*

Three consequences follow, and the third is counterintuitive.

It is a tailwind for automation. The industry's own stated answer — from consultancies, from the vendors, and from the banks' own capital allocation — is "automate, because we cannot hire and retain the expertise." An AI-native approach is not fighting the industry's direction; it is the industry's stated direction. *[INFERENCE]*

It is a threat to any outside vendor. The banks are building their own agentic operations tooling; JPMorgan's cited thirty-to-forty-per-cent reductions are attributed to internal AI. *[FACT]* An outside vendor must integrate with in-house automation rather than compete against it, because the institution has already decided to build. *[INFERENCE]*

And it makes migration harder, not easier. This is the counterintuitive result and it is important. Replacing a core servicing platform is precisely the activity that requires deep platform knowledge — the knowledge that is walking out the door. A bank with a hollowed-out operations function is *less* able to undertake a migration, not more. *[INFERENCE]* The knowledge cliff strengthens Layer 4 of the incumbent's moat (Chapter 30) rather than weakening it.

That third point deserves emphasis because the opposite is usually assumed. The departure of the experts is frequently cited as a reason banks will finally modernise. On this analysis it is a reason they will find modernisation harder and will prefer approaches that do not require it. *[INFERENCE]*

66.4 The knowledge-capture opportunity

One genuine opportunity sits inside this. Capturing retiring practitioners' judgment as durable, queryable, governed knowledge converts a depreciating human asset into a compounding institutional one. *[INFERENCE]* The academic and practitioner literature on AI-assisted knowledge preservation is real and growing. *[FACT]*

The honest caveats: it is not a product category with an established buyer, no institution in this market has been documented purchasing it, and the value is realised over years rather than quarters. *[UNKNOWN]* It is a genuine idea without a demonstrated market.

The honest capability ledger

CATEGORY	CHANGED BY AI?	ASSESSMENT
1 · Calculation	No, and should not	Deterministic systems already do this correctly.
2 · Interpretation	Yes, substantially	Newly economical at scale. This is the real change.
3 · Coordination	Marginally	Consent still requires a counterparty to decide.
4 · Data synchronisation	Partially	Matching automates; investigation needs interpretation and coordination.
5 · Institutional control	No	Who holds the register — untouched.

Pulling Part Nine together against the framework of Chapter 26.

COMPLEXITY CATEGORY	CHANGED BY AI?	ASSESSMENT
1. Calculation	No, and should not	Deterministic systems already do this correctly. AI has no role and introducing one would be a defect. <i>[INFERENCE]</i>
2. Interpretation	Yes, substantially	Reading bespoke agreements and unstructured notices at scale is newly economical. This is the real change. <i>[INFERENCE]</i>
3. Coordination	Marginally	Latency of transmission and status tracking improve. Consent still requires a counterparty to decide. <i>[INFERENCE]</i>
4. Data synchronisation	Partially	Matching and classification automate; investigation requires interpretation and often coordination. <i>[INFERENCE]</i>
5. Institutional control	No	Who holds the register, who has the contractual right to act, who bears liability — untouched. <i>[INFERENCE]</i>

67.1 The verdict

Something real has changed, in exactly one place, and it is not the place where this market's modernization attempts have historically failed. *[INFERENCE]*

Chapter 48 established that every wave that failed, failed in Categories 3 and 5. Chapter 26 established that Category 2 was written off rather than attempted. AI has un-written-off Category 2.

That is genuinely new. Whether it is decisive depends on a question this part cannot answer and Part Eleven must: **is there a strategy in which solving Category 2 for a single institution creates enough value to be bought, without requiring the Category 3 and 5 problems to be solved first?** *[INFERENCE]*

If yes, then AI has changed the answer. If no, then AI has changed the technology and not the market, and the paper's central question has the disappointing answer.

The remainder of the paper is an attempt to determine which.

X

PART TEN

Authority and truth in a multiparty credit market

Everything so far has been description. This part is the paper's one piece of sustained conceptual work, and it addresses the question that all the failed initiatives were unknowingly attempting to answer.

The question is: **in a market where several institutions each maintain their own version of the same facts, what does it mean for one of them to be right?**

That is not a philosophical flourish. It has a precise legal answer in this market, and the answer explains why the distributed-ledger wave failed, why the reconciliation utilities can only detect breaks rather than eliminate them, why the bank consortia are structured as they are, and what the actual boundary of any new entrant's ambition can be.

The preceding nine parts described the market from different directions: legal form, participants, operating lifecycle, incumbent systems, standards, modernization attempts, collaboration, private credit and AI.

This part asks what those observations have in common.

The answer requires separating concepts that market technology discussions routinely collapse into one another: authority, interpretation, calculation, verification, provenance and synchronization.

Once separated, several apparent contradictions in the preceding history become easier to explain.

Who owns the truth about a loan?

Start with the crisp version, then complicate it.

The crisp version: the administrative agent's register is authoritative on ownership. *[FACT]* The agent maintains the register of lender names, commitments and principal amounts, updating it from each assignment and acceptance; a buyer funds upon the agent's notice that a trade is effective in the register; and the agent pays interest on the payment date to the lender of record. *[FACT]*

This is not a convention or a market practice. It is a contractual allocation of authority, written into the credit agreement, and it means something specific: **the register, not the lender's own record, is the operative statement of who owns what.** A lender whose books disagree with the register does not thereby change the register; it has a query, and the burden of pursuing it. *[INFERENCE]*

Two qualifications keep that from being an overstatement, and both matter.

The register is authoritative, not infallible. Credit agreements typically make the agent's records conclusive *absent manifest error*, and the post-Revlon erroneous-payment machinery — under which a notice of erroneous payment is "conclusive absent manifest error" and off-schedule payments are presumptively returnable — exists precisely because agents make mistakes that have to be unwound. *[FACT]* The authority is procedural: it determines who states the position and who bears the burden of challenging it, not who is factually correct.

Authority over ownership is not authority over everything. The register settles who the lenders are. It does not settle what a covenant means, what a position is worth, or whether the collateral exists. §68.1 takes that apart.

With those qualifications, the point stands and is stronger than it is usually put: this is a firmer allocation of authority than exists in most financial markets, and it is why this market's data problem is not, at root, a data problem.

68.1 The complication

The crisp version is true and incomplete, because different facts about a loan have different authorities, and the differences matter operationally.

The borrower is authoritative on its own condition: the financial statements, the compliance certificate, the covenant calculation, the borrowing request, the certification that no default has occurred. The agent and the lenders receive these and may check them, but they do not produce them. *[FACT — ESTABLISHED BY THE CREDIT AGREEMENT'S REPORTING COVENANTS]*

The agent is authoritative on the register, on the calculation of interest and fees under the agreement, on the receipt and distribution of payments, on the effectiveness of assignments, and on the tabulation of votes. *[FACT]*

The lender is authoritative on its own economic position — what it actually holds, what it has agreed to sell, what it has hedged — and on its own accounting treatment, which may legitimately differ from another lender's for the same instrument under different accounting regimes. *[INFERENCE]*

The custodian is authoritative on where the cash and the assets actually are. *[INFERENCE]*

The trustee, in a CLO, is authoritative on the waterfall computation and on the reports that determine investor entitlements. *[INFERENCE]*

The fund administrator and fund accountant are authoritative on the net asset value they strike, which is the number the fund's investors transact on. *[INFERENCE]*

And in private credit, the manager is authoritative on valuation — a fact that has no analogue in the syndicated market and that Chapter 56 showed the supervisors regard as the sector's central vulnerability. *[FACT]*

68.2 The consequence

There is no single golden source, and the reason is not fragmentation or neglect. **The reason is that different questions have different correct answers-holders, and the market has allocated each one deliberately.** *[INFERENCE]*


This is the misdiagnosis at the root of a great deal of failed effort. "The market lacks a golden source" is usually said as though the golden source is missing and should be built. It is not missing. It is **distributed by design**, and the design is contractual.

What the market actually lacks is not a single source of truth. It is a **reliable mechanism for propagating each authority's answer to everyone who needs it, fast enough and with enough evidence that the recipient can rely on it without independently reconstructing it.** *[INFERENCE]*

Those are entirely different problems, and confusing them is why the ledger wave built the wrong thing.

FIGURE 11

No participant sees the whole loan



Different facts are visible to different institutions, at different times, under different authority — and some are visible to nobody. Chapters 68 and 69 take that apart: the market does not lack a source of truth so much as a mechanism for propagating each authority's answer to the parties who need it. *[INFERENCE]*

Authority by data element

The abstract point becomes usable when made concrete. The following table maps the market's principal data elements to their authority, and to what a downstream party can and cannot do about them.

DATA ELEMENT	AUTHORITATIVE PARTY	DOWNSTREAM PARTY'S POSITION
Lender of record, commitment, principal outstanding	Agent (register)	Derivative. Can reconcile; cannot override.
Interest rate for a period	Agent , applying the agreement	Can recompute independently from the agreement and the reference rate; a disagreement is a dispute, not an error to be corrected unilaterally.
Amount of a payment and its allocation	Agent	Derivative. Recomputation is possible and is the basis of every shadow-accounting business.
Effectiveness of an assignment	Agent (plus borrower consent where required)	Cannot be determined by the parties to the trade alone.
Result of an amendment vote	Agent (tabulating), against the register	Derivative.
Covenant compliance	Borrower certifies; agent/lenders may test	Contested territory; the certificate is the record, the calculation may be checked.
Priority of a security interest	The public filing system , by statute	Determinable from the public record. Establishes rank, not that anything exists to rank over. <i>[FACT]</i>
Existence of the pledged asset	Sometimes : the account debtor for a receivable, the title registry for a vehicle	Authoritative in principle, and not consulted in the 2025–26 failures. Chapter 73. <i>[INFERENCE]</i>
Whether an asset has been pledged once	Nobody — no party holds the union of the facilities	The narrow structural gap the recent failures ran through. Chapter 73. <i>[INFERENCE]</i>
Market price / mark	Nobody — vendor services supply quote-derived evaluations	No authority exists. Every holder marks on a purchased or internal estimate. <i>[FACT]</i>
Fair value in private credit	The manager	No external authority exists. This is what the supervisors object to. <i>[FACT]</i>
Net asset value of a fund	Fund administrator / accountant	Investors transact on it.
Accounting treatment	Each holder , under its own regime	Legitimately divergent; not a reconciliation break.

Three observations from the table.

The rows where no authority exists are the rows where the market's recent failures occurred. No party is authoritative on whether a given asset has been pledged once, and the 2025–26 collapses ran through exactly that gap — though Chapter 73 shows that the neighbouring question of whether the asset *exists* usually does have an authoritative answer, which was not sought. *[INFERENCE]* Private-credit fair value has no external authority, and that is the supervisors' named concern. *[FACT]* Market price has no authority, and the market lives with quote-derived estimates.

But the three authority-less rows are not one problem, and treating them as one is an error this paper made in an earlier draft. They differ on the question that determines whether anything can be done about them: *is there a fact of the matter?*

	IS THERE A VERIFIABLE EXTERNAL FACT?	WHAT "VERIFICATION" WOULD MEAN	WHO WOULD BUY IT
Collateral existence and uniqueness	Yes. A receivable either exists or does not; it has been pledged once or more than once.	Establishing the fact against independent evidence	A lender, after a loss or in anticipation of one
Private-credit fair value	No. In the absence of a market price, value is a considered judgment about future cash flows.	Auditing the <i>process</i> — inputs, methodology, consistency, review — not checking the answer	The manager whose judgment is in question, or an LP or supervisor who does not buy software
Market price (BSL)	Partially. Trades occur but are not reported.	Observing what the market actually did	The whole buy side, which cannot obtain it (Chapter 53)

[INFERENCE]

The distinction is decisive later. Collateral verification is a **fact-finding** problem with a loss-prevention buyer. Private-credit valuation is a **process-governance** problem whose most natural buyer is the party being governed — a conflict that has no equivalent in the collateral case. Chapter 73 develops the first; Chapters 81 and 94 keep the two separate for this reason.

The rows where an authority exists are stable, however inconvenient.

Nobody disputes who the lender of record is; they merely find out slowly. *[INFERENCE]*

The distinction between a break and a difference is not made

systematically. A lender whose position differs from the agent's has a break. A lender whose accounting treatment differs from another lender's does not. Systems that treat all differences as exceptions generate noise; systems that understand the authority map generate signal. *[INFERENCE]* This is a genuine and underexploited product insight, and it is the strongest argument this paper can make for why a governance-aware layer would outperform a generic reconciliation engine. *[INFERENCE]*

One golden source, or many with governed synchronization?

The strategic form of the question is whether the market needs one universal golden source or multiple authoritative sources with governed synchronisation.

70.1 The case for one

A single source eliminates reconciliation by construction. If everyone reads the same record, no two records can differ. This is the intuition behind the bank consortium data platform: capture the agent's books and records in real time and centralise them, so lenders read the agent's data rather than reconstructing it. *[FACT]*

The approach has real traction — more than \$4.1 trillion of reported facility coverage — and it is the most successful modernization in the market.

[VENDOR CLAIM]

70.2 The case against one

Coverage will never be complete. The platform's progress is gated by agent onboarding, one bank at a time, and its stated target of roughly ninety per cent of broadly syndicated loans is both ambitious and partial. *[FACT/VENDOR CLAIM]* A lender with a diversified portfolio will hold positions from agents on the platform and agents not on it, and must therefore maintain the reconciliation capability regardless. *[INFERENCE]*

It addresses one authority. The register and the agent's calculations are centralised. Custody, valuation, accounting, collateral and covenant compliance are not, and cannot be, because their authorities are elsewhere. *[INFERENCE]*

Private credit has no equivalent chokepoint. A market where the manager is the agent has no external party whose books can be centralised. *[INFERENCE]*

And a single source concentrates a single point of failure. The 2024 and 2026 security incidents at a vendor whose systems touch wire and transfer instructions for a large share of the world's major banks are a reminder of what concentration means operationally. *[FACT]*

70.3 The judgment

The market will not converge on one golden source, and the evidence for that is thirty years long. *[INFERENCE]* Identifiers did not converge; documentation did not converge; messaging did not converge. The successful modernization is a *partial* centralisation gated by voluntary participation and compensated by ownership.

The realistic architecture is therefore **many authoritative sources with governed synchronisation**: each fact held by its rightful authority, propagated with provenance, and reconciled continuously by parties who understand which differences are breaks and which are legitimate. *[INFERENCE]*

That is a less elegant answer than a single golden source and it is the one the market's structure actually supports.

70.4 What "governed" has to mean

The phrase is doing real work and should not be left as a slogan. Synchronisation is *governed* when five things are specified in advance rather than discovered during a dispute. *[INFERENCE]*

Who is authoritative for this fact. Not "who sent it" — who is entitled to state it. Chapter 69's table is the specification, and it is a contractual question with a contractual answer, different for each row.

What event triggers propagation. A fact changes on an event: an assignment becomes effective, a rate is set, a payment is made, an amendment passes. Propagation on an event is a different architecture from propagation on a schedule, and the market's current default — a notice, sent when someone gets to it — is neither.

How stale the recipient's copy is permitted to be. Latency is a governed parameter, not an accident. The difference between a position known at the moment of assignment and one known at the next reconciliation run is the difference between a control and a report.


What evidence travels with the fact. Source, method, version, reviewer — the content of Chapter 71.1. A fact that arrives without its evidence cannot be relied on without being reconstructed, which is the reconciliation cost the whole exercise is trying to remove.

How a difference is classified when one appears. This is the part the market does worst. A downstream record that differs from the authority's is a **break**. A downstream record that differs because it is kept under a different accounting regime is a **difference**, and is correct. A system that cannot tell them apart produces exception queues that operations teams learn to ignore, which is worse than producing none. *[INFERENCE]*

Governed synchronisation is therefore not a technology. It is an agreement about those five things, which some technology then enforces. That distinction is why the ledger wave built the wrong artefact, and Chapter 71 is the reason.

FIGURE 12

Authority is allocated fact by fact



Different facts have different rightful authorities; several have more than one legitimate source; and some have none, the absence being itself a material finding. Chapter 70 declines the single golden source as the target architecture. Case labels preserve the status distinctions used in Chapter 73 and Appendix F. [FACT]

Provenance versus consensus

This chapter explains, in one distinction, why the distributed-ledger wave failed.

Consensus answers: how do mutually distrustful parties agree on a shared state without a trusted intermediary? That is a real problem with real solutions, and distributed ledger technology is one of them.

Provenance answers: how does a party receiving a fact know where it came from, who asserted it, when, on what evidence, and whether it has changed since?

The syndicated loan market's problem is the second, not the first. *[INFERENCE]*

The parties are not mutually distrustful peers. They are a principal and counterparties in a hierarchy of authority established by a negotiated contract. Nobody disputes that the agent decides; they merely cannot see what the agent decided, in time, with confidence. There is a trusted intermediary — the agent — and it is trusted contractually whether or not it is trusted personally. *[INFERENCE]*

The independent evidence supports this reading directly. The barriers to DLT adoption in loans were legal and regulatory uncertainty, network interoperability and fragmentation, scalability, and the finding that **for private permissioned networks a centralized system is often faster, cheaper and better.** *[FACT]* The binding constraint was coordination, legal enablement and network effects among agent banks — not compute. *[FACT]*

A permissioned ledger among parties who already have a contractually designated authority is a distributed database with extra steps. *[INFERENCE]* It supplies consensus where consensus was not the missing ingredient, and it does not supply provenance, which was.

71.1 What provenance actually requires

If the market's need is provenance, it is worth specifying what that means, because the word is used loosely.

For a given assertion about a loan — a margin grid, a position, a payment amount, a covenant level — a provenance-bearing system must be able to state:

- **The source:** which document, which clause, which notice, which counterparty transmission.
- **The assertion:** what was extracted or computed, in a structured form.
- **The method:** whether a human read it, a model proposed it, a deterministic calculation produced it.
- **The confidence and the reviewer:** what was checked, by whom, when.
- **The version:** which amendment of the agreement, which version of the register, as of what date.
- **The history:** every prior value and the reason it changed.

None of that requires a ledger. All of it requires discipline, and the discipline is the product. *[INFERENCE]*

This is also precisely what a data incumbent has begun marketing under the name "source-cell traceability" in private credit. *[FACT]* The concept is not proprietary and the position is contested.

71.2 Three words that are routinely confused

The argument of this part depends on keeping three things apart, and most writing about this market runs them together. *[INFERENCE]*

	WHAT IT IS	WHO SUPPLIES IT	WHAT IT CANNOT DO
Authority	The contractual or legal entitlement to state a fact	Designated in advance — the agreement, the statute	Guarantee that the stated fact is correct
Provenance	A property of a record: where the assertion came from, who made it, on what evidence, as of when	The system that carries the fact	Make the assertion true, or confer the right to make it
Verification	An act performed on a fact: independently establishing it against evidence the asserting party does not control	Anyone with access to the evidence	Confer authority, or create evidence that does not exist

The market's failures divide cleanly along these lines. Where authority exists and provenance does not, the result is reconciliation — everyone reconstructing what the authority already knows (Parts Two and Three). Where authority exists and someone verifies anyway, the result is evidence and earlier detection, which is worth money but is not control (Chapter 72). And where no authority exists at all, verification is the *only* available instrument, which is why the collateral case in Chapter 73 is both the strongest and the hardest.

Chapter 72 takes the second case. Chapter 73 takes the third.

What verification can mean when the agent is contractually the truth

A hard question that any verification-oriented thesis must answer honestly.

If the agent's register is authoritative by contract, what is the value of an independent system that recomputes it?

The answer is not "catching the agent's errors and overriding them," because the system cannot override the agent. The answer has four parts, and they are worth separating because only some of them are strong.

One: earlier detection. A lender that independently computes what its interest payment should be knows on the payment date whether the received amount matches. Today it often learns later, via a reconciliation run against data derived from the same source. *[INFERENCE]* Earlier detection is worth something in working capital and in dispute leverage.

Two: evidence for the dispute. When a lender and an agent disagree, the party with a reconstruction traceable to the credit agreement's clauses has a materially stronger position than the party asserting that its spreadsheet says otherwise. *[INFERENCE]* Post-Revlon, the market has erroneous-payment machinery in which such evidence has direct legal consequence. *[FACT]*

Three: internal governance. A lender's own audit, risk and compliance functions need to demonstrate that they have controls over positions and cash. An independent recomputation is a control. *[INFERENCE]* This is a genuine and underrated buyer: the purchaser is not the operations team seeking efficiency but the control function seeking evidence.

Four: the collateral case, which is different and stronger. Verification of the register is verification against a party with authority. Verification of *collateral* is verification in a domain where the authority is allocated

differently — and where, for asset classes with no per-asset record, the narrow question of whether an asset has been pledged once is one no party holds the answer to. *[INFERENCE]*

That claim has to be stated carefully, because a lawyer will otherwise supply the correction, forcefully. **Public security-interest filing regimes exist and are precisely designed to give notice of competing claims**, and they do that job. A financing statement puts the world on notice that a secured party claims an interest in a category of collateral, and the statutory priority rules then rank the claims deterministically. To say that "nobody is authoritative about whether an asset was pledged twice" would be wrong as a general statement about secured lending.

The claim this paper makes is narrower and is set out in full in the next chapter: **the filing system establishes rank rather than existence, and for asset classes with no per-asset record — trade receivables above all — no party holds the fact of whether a given asset has been financed more than once**. Chapter 73 takes it apart against the statutory text and the recent cases, and Appendix F carries the mechanics.

72.1 The honest limit

Verification produces **information and evidence**. It does not produce **authority**. *[INFERENCE]*

A system that recomputes a loan perfectly and disagrees with the agent has not changed what the loan is. It has given its owner a well-founded reason to raise a query, and a stronger position in the resulting conversation.

That is worth money. It is not worth what a system of record is worth. Any business plan that prices verification as though it were authority is mispriced, and any strategy that assumes verification is a path to authority needs to specify the mechanism by which the transition occurs.

The fourth of those four parts is the one that can be tested against evidence rather than argued from structure, and the next chapter does that. The chapter after it returns to the question of access.

Collateral: when legal authority does not establish economic fact

The preceding five chapters argued that this market's difficulty is not a missing truth but an unpropagated authority, and that the places where it fails hardest are the places where no party is authoritative at all. That argument has so far been made from the structure of credit agreements. Collateral is where it can be tested against something else, because in 2025 and 2026 the market ran the experiment, at named institutions, for documented amounts.

Collateral is also the hardest available test, and that is the reason to run it. Of all the domains this paper has examined, collateral is the one where a comprehensive, mandatory, public legal infrastructure already exists and has existed for decades. If the paper's framework survives contact with a domain that is *already* legally instrumented, it is worth something. If it does not, Chapter 69's authority table is a description of a market that simply has not got around to building things.

73.1 Article 9 is not the failure

The first thing to say is the thing most commentary gets wrong.

UCC Article 9 is a **notice-filing** system. Its Official Comment says so in terms: what is filed is not the security agreement but "a simple record providing a limited amount of information," and "the notice itself indicates merely that a person **may** have a security interest in the collateral indicated. Further inquiry from the parties concerned will be necessary to disclose the complete state of affairs." [FACT — UCC § 9-502, OFFICIAL COMMENT 2]

A financing statement requires three things: the debtor's name, the secured party's name, and an indication of the collateral. [FACT — § 9-502(A)] The filing office's role is ministerial: the grounds on which it may refuse a record are exhaustively enumerated and are all formal — wrong medium, unpaid fee,

missing name or address. None of them permits the office to ask whether the collateral exists. The Official Comment states that a filing office may not reject a record even where the information it contains is facially implausible. *[FACT – § 9-516 AND OFFICIAL COMMENT]*

This is not a defect. It is the design, and the design works. A searching creditor can establish that a prior claim is asserted, can date it, and can identify who asserts it, which is what the system was built to deliver.

Article 9 allocates legal rights, and it allocates them successfully. The 2025–26 failures were not failures of that allocation. *[INFERENCE]*

The correct statement of the finding is therefore narrower and, stated properly, more interesting: **a legal system can allocate rights completely and correctly without establishing every underlying economic fact to which those rights relate.**

73.2 Filing is not attachment; priority is not integrity

Four concepts are routinely collapsed in commentary about these events, and keeping them apart is what makes the argument precise rather than rhetorical.

Filing is the public act. § 9-502(d) expressly permits a financing statement to be filed *before* a security agreement is made or any security interest attaches. *[FACT]*

Attachment is what makes a security interest enforceable, and it requires three things that appear nowhere on the filing: value given, an authenticated security agreement describing the collateral, and — decisively — that **the debtor has rights in the collateral.** *[FACT – § 9-203(B)]*

Perfection is attachment *plus* the applicable perfection step. *[FACT – § 9-308(A)]*
A filing without attachment does not produce a perfected interest.

Priority ranks perfected claims, ordinarily from the earlier of first filing or first perfection. *[FACT – § 9-322(A)(1)]*

The consequence for a fabricated receivable is exact, and it is not the one usually stated. It is not that a perfectly perfected security interest existed over collateral that did not exist. A debtor has no rights in an asset that does not exist, so **attachment fails** under § 9-203(b)(2), and an interest that has not attached cannot be perfected. What remains is this: **a facially valid public filing can exist without establishing that an enforceable security interest ever attached to the represented economic asset.** *[INFERENCE – FROM THE STATUTORY TEXT]* The paper carries that formulation and not the looser one.

There is a related asymmetry worth recording because it is the mechanical reason the public record is coarse. A financing statement may describe collateral supergenerically – "all assets," "all personal property." *[FACT – § 9-504(2)]* A security agreement may **not**: § 9-108(c) provides that such a description does not reasonably identify collateral. *[FACT]* The public layer is deliberately permitted to be vaguer than the private one. A searcher therefore learns that someone claims accounts; the question of *which* accounts is answered only in documents the searcher cannot see.

Two further provisions close the granularity question. Under § 9-204(a) a security agreement may cover after-acquired collateral, so an agreement made today reaches invoices that do not yet exist – the interest attaching to each as the debtor acquires rights in it, which is precisely why none of them can be enumerated in the filing. *[FACT]* And the statutory right to demand an accounting of the collateral and the obligation runs to the **debtor only** – Official Comment 3 to § 9-210 says a third party is not entitled to request that information. *[FACT]* The public system is not merely silent at the individual-asset level; the private system's information rights are pointed away from the party who would most want to ask.

Priority is a legal ranking. It does not establish that anything is there to rank over. *[INFERENCE]*

73.3 The asset class decides — and Tricolor decides against the easy version

The argument would be simpler if no per-asset registry existed anywhere. One does, and the most instructive of the recent failures happened on top of it.

Tricolor's collateral was subprime auto loans. A vehicle securing such a loan has a state title record, and in the states that operate electronic lien-and-title programmes those systems maintain an authoritative, per-vehicle record of the noted lienholder. *[FACT]* That is closer to asset-level uniqueness infrastructure than anything receivables have. Bloomberg, reporting from the bankruptcy record, put the number of active loans sharing identical vehicle identification numbers across facilities at approximately 29,000 of roughly 70,000. *[REPORTED]* The SEC alleges that on one occasion 3,225 identical loans were directed into two separate securitizations. *[FACT — THAT THIS IS THE ALLEGATION; SEC PRESS RELEASE 2026-77, NOT ADJUDICATED]* Warehouse lenders worked from loan tapes supplied by the borrower. A cross-facility comparison of vehicle identification numbers, and a query of independent title records, were both available and neither was performed. *[FACT — AS REPORTED IN THE CASE ANALYSES]* The reported assessment that either would have exposed the duplication rests on trade and legal-press analysis rather than on a court finding. *[FACT — AS REPORTED]*

The verification mechanism existed, was authoritative, was operational, and was not used. *[INFERENCE]*

That single observation does more work than any other in this chapter, because it disciplines the conclusion. Where an authoritative per-asset record exists, the residual gap is **operational**, not infrastructural, and the honest description of what failed is that nobody queried it.

Petters makes the same point from further back. Purchase orders purportedly issued by two named national retailers were fabricated; the retailers themselves were authoritative and could have confirmed or denied every one of them; the note purchaser's manager did not ask. The

scheme ran to \$3.65 billion and its principal was convicted and sentenced to fifty years. *[FACT – CONVICTION AND SENTENCE]* No infrastructure was missing. A telephone call was missing. *[INFERENCE]*

73.4 Receivables are where the boundary actually sits

First Brands is the case that cannot be disposed of this way, and it is therefore the one that locates the real boundary.

Trade receivables have no equivalent of a vehicle title. There is no public per-invoice record, no authoritative party who can be queried about whether a given invoice – number 41,772, say – exists and has been financed once, and – by the architecture set out in §73.2 – no prospect of the filing system supplying one. Multiple factoring counterparties held filings and security agreements against overlapping receivables pools. Invoices were fabricated or inflated; the same invoices were reported to more than one counterparty; field examinations received aggregated schedules rather than original documents; an audit that surfaced discrepancies was explained away; and information requests were routed through the people conducting the scheme. The reported quantum of fabricated or inflated receivables is approximately \$2.3–2.7 billion depending on which measure is cited, and the principals have been criminally charged. *[FACT – CHARGES; THE QUANTUM IS AS REPORTED, AND APPENDIX F SETS OUT WHICH FIGURE MEASURES WHAT]*

The authoritative source for a trade receivable is the **account debtor** – the customer who either owes the money or does not. That party is reachable: § 9-406(a) contemplates notification to the account debtor, and independent confirmation directly from the obligor is a recognised category of high-reliability audit evidence. *[FACT]* It was available at First Brands and was not used at scale. *[FACT – AS REPORTED]*

So even here, part of the gap is operational. What is *not* operational is the second half of the question. Confirming with the customer that an invoice is genuine and unpaid does not tell a lender whether the same invoice is

sitting in another lender's borrowing base. That fact exists in no public record, in no authoritative party's hands, and in no document any single lender is entitled to demand. *[INFERENCE]*

73.5 The strongest counterargument, stated at full strength

A secured-lending lawyer and an asset-based-lending credit officer will both make the same objection, and it is a good one.

The objection is that these were underwriting and audit failures dressed up as an infrastructure story. Independent collateral verification was once ordinary market practice and was competed away. Red flags were present in every case and were not acted on. Field examinations are prescribed in supervisory handbooks for asset-based lending. Unannounced audits, reconciliation of borrower data tapes against independent servicing records, cash dominion and lockbox arrangements, third-party custody of chattel paper, and direct obligor confirmation are long-established practice rather than supervisory requirement. All of them were available to every lender that chose to use them. On this reading the market does not need new infrastructure; it needs the discipline it used to have.

This objection is substantially correct, and the paper concedes it.

[INFERENCE] In each of the three cases above, better single-institution diligence would have reduced or prevented the loss. Any account of these events that presents them as purely structural is overstating, and this paper's earlier draft came close to doing so.

73.6 Where individual diligence stops

The objection is nonetheless incomplete, and the place where it stops is precise.

Diligence is a capability of one institution applied to its own exposure. Uniqueness is a property of the relationship *between* institutions. A lender performing flawless verification of its own borrowing base still cannot see that the same receivable has been pledged into a facility at a different

funder, because no lender can see another lender's facility, no lender can compel disclosure of it, and — where no per-asset registry exists — no party holds the union of the two. Where such a registry does exist, as it does for titled goods, the constraint does not bind. *[INFERENCE]*

That is the structural residue, and it is small, specific and real. The paper claims nothing wider. Existence, eligibility and the accuracy of borrower reporting are addressable by an institution acting alone, at a cost it can choose to bear; valuation is a different problem and is treated separately (§87.2). **Uniqueness across facilities, in the asset classes that have no per-asset record, is not.** *[INFERENCE]*

73.7 The infrastructure question is not automatically a registry

Having identified a gap, the reflex is to specify a registry. The reflex should be resisted, because at least four different mechanisms could address parts of the gap and they have very different coordination requirements.

Query the authoritative record where one already exists. For titled assets this is available today and is the Tricolor lesson: no new institution required, no counterparty's consent required, one lender's decision.

Confirm with the party who actually knows. For receivables the account debtor is authoritative on the underlying obligation. Confirmation at portfolio scale is an operational and evidential problem — how to obtain it, how to keep it current, how to prove it later — not a coordination problem.

Move the asset into custody or control. Third-party custody of the underlying paper, and cash dominion over the collections, substitute possession for representation. These are contractual and bilateral.

Match across institutions without disclosing. Privacy-preserving comparison of asset fingerprints across lenders is the only one of the four that requires other institutions to participate — and it is the only one that can answer the uniqueness question.

Three of the four escape Chapter 41's mechanism, because the party bearing the cost is the party receiving the benefit and **no other lender** has to agree to anything. Two of those three still need a counterparty's cooperation — the account debtor must answer, the borrower and a custodian must contract — but that is a bilateral negotiation, not an industry one. *[INFERENCE]* The fourth requires the other lenders, and §73.8 is about why that matters.

There is also an intermediate class worth naming precisely, because it is currently marketed as more than it is. Digitising a document, hashing it, or recording it on a shared ledger establishes that a document matching a previously registered fingerprint has been presented before and has not been altered since it was registered. **A hash of a fabricated invoice is still evidence of a fabricated invoice.** *[INFERENCE]* Document existence is not economic-asset existence, and the two are routinely conflated in materials describing this problem. The distinction of Chapter 71 applies unchanged: these are provenance mechanisms, and provenance is genuinely useful, but it is not verification of the underlying fact.

73.8 The coordination trap

The cross-institution mechanism has a property that determines whether it can be built at all, and it is the property this paper has spent Part Five and Part Seven describing in other settings.

A shared uniqueness mechanism is non-linear in adoption. A network covering most lenders catches most duplicate pledges and misses entirely any pledge routed through a lender outside it. Partial coverage therefore delivers a **warning signal**; only complete coverage of the relevant market delivers a **uniqueness proof**. One non-participant is sufficient to convert the second into the first. *[INFERENCE]*

The precedent is the filing system itself. Notice filing achieved near-universal coverage for the collateral it governs — accounts among them — because filing was made the default requirement for perfection and

therefore the route to priority, not because the market agreed it was a good idea. *[FACT – THE STATUTORY STRUCTURE]* Whether an equivalent uniqueness obligation could be imposed by supervisory action rather than by amending the Code is a question for counsel and is listed in Appendix F. *[UNKNOWN]*

Two consequences follow. First, the parts of this problem that a single institution can solve alone are available now, and the parts that require an industry are subject to the same thirty-year base rate as everything else in this paper. Second, a warning signal is not worthless — it is simply a different and smaller product than a proof, and it should be priced and described as one. *[INFERENCE]*

It should also be recorded that this is a moving target. Infrastructure aimed at parts of the gap is under active construction — a tokenised-collateral platform announced by a central clearing utility for centrally cleared securities, the digital-asset amendments to the Code enacted in a majority but not all states, and vendor products marketing cross-facility duplicate-pledge detection. *[FACT / VENDOR CLAIM]* None of it covers the asset classes and market segments where the recent losses occurred. *[INFERENCE]* The gap is being addressed. It is not closed, and a claim that it is closing on a stated timetable is not supportable from the public record. *[UNKNOWN]*

73.9 What this establishes

Return to the authority table.

Chapter 69's table now carries three collateral rows where an earlier draft of this paper carried one. That split is the finding of this chapter, and it is set out here with what the record supports for each:

QUESTION	IS THERE AN AUTHORITY?	WHAT THE 2025–26 RECORD SHOWS
Who ranks ahead of whom	Yes — the public filing system, by statute	Worked as designed; not implicated in the failures
Does the asset exist	Sometimes — the account debtor for a receivable; the title registry for a vehicle	Authoritative sources existed in every case examined and were not consulted
Has it been pledged once	For titled goods, yes — the title record. For receivables, no — no party holds the union of the facilities	The genuine structural gap, and it is confined to the asset classes with no per-asset record

[INFERENCE]

The general proposition the chapter establishes is the one stated at the outset, and it now has evidence behind it rather than reasoning alone: **legal authority, economic fact, interpretation and verification are related but different things, and a market can have the first in excellent order while the second goes unestablished for years.** [INFERENCE]

This is the same structure the paper has met repeatedly under other names. The register is authoritative on ownership and silent on value. The agreement is authoritative on the calculation and silent on whether the borrower's certificate is true. The filing is authoritative on rank and silent on existence. In each case the authority is real, the allocation is deliberate, and the gap is not where the institutions are looking.

73.10 The implication, stated carefully

The implication is not that Article 9 should be replaced. Nothing in the record supports that, the system is doing what it was built to do, and a proposal to rebuild the law of secured transactions in order to solve an operational verification problem would be the same category error the ledger wave made in Chapter 71. [INFERENCE]

The implication is narrower: **the available work is governed verification of economic facts that the existing legal and operational infrastructure does not continuously establish** — with the emphasis on *continuously*, because

the mechanisms that failed were periodic, sample-based, and dependent on the borrower for their inputs. *[INFERENCE]*

One boundary belongs here rather than in the appendix, because it is the place readers most often over-claim. Machine reading genuinely shortens this work: it can read every document rather than a sample, structure them, resolve obligors across facilities, and surface the contradictions that should prompt a confirmation nobody currently makes. What it cannot do is establish the fact. **AI can detect a reason to doubt; it cannot manufacture authoritative evidence that does not exist** — and where the authoritative record is absent, as it is for the individual receivable, no analytical technique creates one. *[INFERENCE]* Appendix F §P sets out the pipeline and the limit; Chapter 65's deterministic boundary is unchanged by anything in this chapter.

Whether that work supports a business is a separate question with a separate answer, and this chapter does not attempt it. A demonstrated problem is not a demonstrated business. Chapter 94 tests it against buyers, budgets and incumbents; Appendix F carries the legal mechanics, the case detail, the matrices and the questions that remain open for counsel.

FIGURE 13

Legal certainty and economic fact are different questions



A facially valid filing establishes rank; it does not establish that the represented asset exists, is collectible, or has been pledged once. Chapter 73 and Appendix F §D separate those three questions. Case labels preserve the quantity classes — ALLEGED, CHARGED (criminal), CHARGED (civil), FOUND — which are not evidence labels. [FACT]

The legal gate

The final and most underestimated constraint. It is not technical, it is not commercial, and it is the one that most often kills approaches of this class.

74.1 The problem

Any system that observes a bank's loan and payment traffic is handling **non-public personal information** and wire instructions. Under the applicable US financial-privacy framework, a financial institution must ensure the security, confidentiality and integrity of that information and remains responsible for ensuring that its service providers and third-party affiliates take the necessary steps to safeguard it. *[FACT]*

A third party inserting a listener onto that traffic becomes an in-scope service provider. **The bank inherits liability for the interception.**

[INFERENCE] Which means the decision is not made by the operations team that wants the capability; it is made by security, vendor risk, and the general counsel, and their default answer to a novel interception is **no.** *[INFERENCE]*

74.2 The precedent, and it is not encouraging

The closest historical analogue to reverse-engineering a system's behaviour from its observed traffic is **screen scraping** in consumer banking. It worked at scale for a decade and is now in regulatory and industry retreat.

The criticisms are documented and they transfer directly: it is "a hack, a workaround... at the mercy of the bank's HTML," it is "fragile," and it carries "privacy, security, and liability risks." *[VENDOR CLAIM — FROM FIRMS SELLING THE SANCTIONED ALTERNATIVE, AND THEREFORE INTERESTED, BUT THE REGULATORY DIRECTION IS INDEPENDENTLY ESTABLISHED]* The European payment-services regime from 2018 and the US open-banking rule finalised in October–November 2024 explicitly push firms off interception and onto sanctioned application programming interfaces. *[FACT]*

The trajectory is unambiguous: **interception strategies get regulated toward consented, auditable, revocable interfaces over time.** *[INFERENCE]*

An entrant designing an interception-based approach should assume it is designing for a transitional regime and should build for consent and auditability from day one, because that is where the regime ends up.

74.3 The academic limit

There is a research literature on this exact problem. Protocol reverse engineering "infers the format of an unknown protocol by analyzing traffic traces," splitting messages into tokens and inferring both message formats and protocol state machines. *[FACT]*

Its key documented limitations bear directly on any traffic-observation strategy: **network-trace-only inference is less accurate than execution-informed inference, and encrypted protocols largely defeat network-only approaches.** *[FACT]*

That second point is a strategic vulnerability, not merely a technical one. An incumbent that wishes to degrade an observer can do so by encrypting, versioning or restructuring its message formats — exactly as website changes broke scrapers. *[INFERENCE]* **An interception surface is only as stable as a party that has every incentive to change it.**

74.4 The three load-bearing risks

Any observe-then-take-over architecture must confront three problems that are usually elided. They are stated here because Part Eleven will test the hypothesis against them.

Risk one: passive observation and active takeover are different asks, and the plan must name where it crosses the line. *[INFERENCE]* The canonical modernization guidance for the pattern states explicitly that it is unsuitable when "requests to the back-end system can't be intercepted" and when "you can't access the legacy system's source code... [because] to disable migrated features and redirect internal calls, you need to be able to modify the

legacy system's source code." *[VENDOR CLAIM — FROM THE PLATFORM VENDOR PUBLISHING THE PATTERN GUIDANCE]* A purely passive listener can observe and shadow indefinitely. The moment it takes over a function, it must insert a **control point in the write path** — a façade that can suppress the legacy system and assert new authority. Mediating live loan and payment instructions is contractually and operationally a heavier ask than observing them, by an order of magnitude. The elegant "just listen" story and the "take over" story are in tension.

Risk two: messages are events, not full state. *[INFERENCE]* Event interception captures *updates to state* — deltas. Any state the servicing platform computes internally and never emits — derived balances, intermediate accrual state, configuration-driven business rules, historical positions predating the listener — is invisible to a traffic observer. *[FACT — THIS IS THE DOCUMENTED CHARACTER OF THE PATTERN]* Worse, even purpose-built event streams drift: delivery is at-least-once rather than exactly-once, meaning duplicates, and on failure an event can be lost with downstream consumers processing a phantom. *[FACT]* Reconstructing an authoritative record from an observed message stream is strictly harder than the outbox problem the software industry already struggles with. Continuous reconciliation is not a nice-to-have; it is the only thing that makes the approach sound.

Risk three: the legal right to intercept is the true gate. *[INFERENCE]*
Everything above.

74.5 The competitive wrinkle

One further fact belongs here because it is frequently missed. **The incumbent already sells the translation layer.** *[FACT]* The vendor ships interface products with open interfaces and a market-integration layer, in which — per its own materials — the vendor manages the translation and integration to the core. *[VENDOR CLAIM]*

The consequence is direct: **the differentiator for any entrant cannot be "we translate the incumbent's data," because the incumbent monetises that translation and can price or bundle against it.** *[INFERENCE]* Whatever the differentiator is, it must be something the incumbent's interface layer does not do and would not want to do.

The model this part has built

Seven chapters is a long way to travel for one architecture, and it is worth assembling before Part Eleven puts weight on it. The pieces are not five adjacent observations. They are one model with five joints, and each joint answers a question the previous one raises.

Start with the question. In a market where several institutions each keep their own version of the same facts, what does it mean for one of them to be right? (Chapter 68)

Answer one: authority is allocated, fact by fact, in advance. The agent is authoritative on the register; the borrower on its own condition; the trustee on the waterfall; the manager, in private credit, on value. This is not fragmentation. It is a contractual design, and it is why "the market lacks a golden source" is a misdiagnosis. (Chapter 69)

Answer two: authority is not infallibility. The register governs operationally, absent manifest error, until it is challenged and corrected. It determines who states the position and who carries the burden of disagreeing — not who is factually right. Every claim in this paper about authority is a claim about procedure. (Chapter 68)

Consequence one: the architecture is many authoritative sources with governed synchronisation, not one source of truth — and *governed* means five specified things: who is entitled to state the fact, what event moves it, how stale a copy may be, what evidence travels with it, and how a difference is classified when one appears. (Chapter 70)

Consequence two: what the market lacks is provenance, not consensus. The parties are not mutually distrustful peers needing to agree without an intermediary. They are a hierarchy with a designated authority, needing to *see what the authority decided*, in time, with enough evidence to rely on it. Distributed ledgers answered the first question. (Chapter 71)

Consequence three: verification is a third thing again — an act, not a property and not an entitlement. It produces information and evidence. Against a party that holds authority, that is worth earlier detection, dispute leverage and a control the audit function will accept. It is not worth what authority is worth, and pricing it as though it were is the most common commercial error in this market. (Chapter 72)


And the limit case: where no party is authoritative, verification is the only instrument there is. Collateral is that case, and it is the hardest test the model faces, because the domain already has a comprehensive public legal infrastructure. The finding is that the infrastructure works — and that it allocates *rights* without establishing the *economic facts* to which those rights attach. Legal authority, economic fact, interpretation and verification are related and different things. (Chapter 73)

One practical rider: any of this that requires observing another institution's traffic runs into a gate that is legal rather than technical, and the trajectory of that gate is toward consent and auditability. (Chapter 74)

That is the model. Part Eleven tests a business hypothesis against it; Part Twelve asks what an industry that accepted it would do differently; Part Thirteen asks what it permits a single firm to build.

FIGURE 14

Many authoritative sources, governed synchronization



The architecture the paper arrives at connects authoritative sources without taking ownership of the truth: identities mapped, provenance preserved, differences reconciled without overwriting the source, and gaps acknowledged rather than filled. The Part Ten closing states the same model in prose. [INFERENCE]

XI

PART ELEVEN

Testing the hypothesis

The preceding ten parts described a market. This part tests a proposition about it.

The proposition under test is the one that most sophisticated observers of this market currently converge on, and it deserves to be stated in its strongest form before it is attacked. It is treated here as a hypothesis rather than a conclusion, in accordance with the paper's method: the market analysis is offered with confidence, the strategic analysis is offered as an argument with exposed premises.

The reader should be warned that this part concludes with a qualified negative on the strongest version of the hypothesis and a qualified positive on a narrower one. Readers looking for validation will not find it, and readers looking for a clean dismissal will not find that either.

The coexistence hypothesis, stated falsifiably

75.1 The hypothesis

Do not attempt to replace the incumbent servicing core. Sit alongside it: observe its message and notice traffic; use AI to read the bespoke agreements and reverse-engineer the message flow into a documented, structured model; run a parallel shadow system that reconciles against the legacy on every position and payment; capture retiring practitioners' judgment as governed knowledge while doing so; then take over functions progressively until the core is isolated and replacing it becomes optional. Aim first at the private-credit and collateral-governance greenfield, where the bank consortium wall is not yet built. Land a single anchor customer fast, before it is.

75.2 The architectural precedent

The hypothesis is not architecturally novel, which is a strength and a warning. It is a composition of five named, published, battle-tested modernization patterns, each with authorship and a known risk profile. *[FACT]*

- The **strangler fig** as the umbrella: a gradual process of modernization that "begins with small additions, often new features, that are built on top of, yet separate to the legacy code base," with a façade intercepting requests and routing each incrementally from legacy to new until the legacy system can be decommissioned. *[FACT]*
- **Event interception** as the passive-listener mechanism: "intercept any updates to system state and route some of them to a new component" — alongside its siblings *legacy mimic* and *divert the flow*. *[FACT]*
- The **anti-corruption layer** as the translation boundary: a façade that translates communication so a legacy system's "convoluted data schemas or obsolete APIs" do not corrupt a cleanly designed modern

application. *[FACT]*

- **Change data capture and event streaming** as the industrial implementation, capturing changes from a source without modifying it and streaming them to a bus that buffers between legacy and new. *[VENDOR CLAIM — THE PATTERN IS DOCUMENTED BY THE PLATFORM VENDORS]*
- **Branch by abstraction and parallel run** as the safe-takeover mechanism, letting old and new implementations coexist behind an abstraction and comparing their results with only one treated as the source of truth. *[FACT]*

That the hypothesis maps cleanly onto named patterns is genuinely useful: it is defensible to a bank's architecture review board, which is a real gate. *[INFERENCE]*

75.3 The bank precedents

Interception-based modernization has been done in banks. Reported cases include a major North American bank using a streaming platform to "rescue data off of its mainframe with a cloud-native, microservice-based approach"; a US regional bank combining change data capture with streaming so that "99.99% of mainframe changes are reflected in cloud applications within four seconds"; and a Danish financial-technology cooperative streaming a forty-year-old core banking mainframe serving more than twenty banks into a customer platform. *[VENDOR CLAIM — THESE ARE THE STREAMING PLATFORM VENDOR'S CASE STUDIES AND SHOULD BE READ AS SUCH]*

Independent press reports a global bank cutting costs ninety per cent modernising its mainframe with stream processing. *[FACT]*

The "sidecar core in parallel" is now described as the dominant recommended model for mid-market bank modernization: a sidecar core runs "in parallel with an existing legacy core, handling a defined subset of customers, products, or services," with an abstraction layer built first, explicitly to reach "eventual core replacement without the associated risks and costs." *[VENDOR CLAIM / CONSULTANCY]*

75.4 The falsifiable form

The hypothesis is composed of six claims, each of which can be independently true or false. Stating them separately is the only way to test it, because the hypothesis is usually defended as a whole and fails in parts.

H1 (Interpretation). AI can read bespoke credit agreements and unstructured notices well enough to produce a structured model of a facility that is accurate enough to be worth verifying.

H2 (Shadow). A parallel system can reconcile against the incumbent on positions and payments accurately and continuously enough to be trusted by the institution's own control functions.

H3 (Access). An institution will grant a third party the access required to observe the traffic, and the legal and security review will permit it.

H4 (Value). The resulting capability is worth enough to an individual institution to be bought, at a price supporting a business, without requiring market-wide adoption.

H5 (Takeover). Progressive takeover of functions is achievable, meaning the institution will eventually let the shadow system mediate live instructions.

H6 (Timing). The private-credit and collateral-governance segment remains unclaimed long enough to establish a position.

The case for

Assembled at its strongest, from the evidence in Parts One through Ten.

The problem is real, large, and independently attested. Roughly seven million agent notices a year moving by fax and email. *[FACT]* Three-week settlement against T+1 for equities. *[FACT]* Single-digit electronic penetration against forty-three per cent in corporate bonds. *[FACT]* A duplication census in which the same interest payment is independently recomputed by eight or more institutions (Chapter 25). Documented buy-side inability "to obtain a comprehensive, real-time view." *[VENDOR CLAIM, CORROBORATED BY THE MARKET'S OWN TRADE-BODY EXECUTIVES]* *[FACT]*

The costs are documented, attributable and recent. A ~\$894 million wire error requiring two years of litigation to reverse. *[FACT]* A separate \$400 million enforcement action against the same institution two months later, for enterprise-wide risk-management, data-governance and internal-control deficiencies — an action the error is widely cited as emblematic of, though the public record does not establish that the error caused the penalty. *[FACT / INFERENCE AS TO THE CONNECTION]* An \$81 trillion near-miss with the same maker-checker failure mode four years later. *[FACT]* And the 2025–26 collateral frauds, which reached named institutions at scale: approximately \$715 million of gross exposure at one, more than \$500 million at another, a \$170 million charge-off at a third. *[FACT — EXPOSURES AND CHARGE-OFFS, NOT AN AGGREGATED LOSS; SEE §59.1]*

A supervisory tailwind now exists. The Financial Stability Board, the Bank of England and the Securities and Exchange Commission have all, in writing, named valuation opacity and data gaps in private credit. *[FACT]* This converts governance from an efficiency argument into a compliance argument, and compliance budgets are less discretionary. *[INFERENCE]*

The human capability is hollowing out and the industry's own answer is automation. One-to-two-year training against twenty-to-seventy-per-cent attrition *[VENDOR CLAIM]*; 2026 back-office layoffs attributed to AI at scale *[FACT]*; and consultancy and vendor consensus that the answer is to automate because the expertise cannot be hired. *[VENDOR CLAIM]* An AI-native approach is the industry's stated direction. *[INFERENCE]*

AI changed the binding constraint. Category 2 — reading the bespoke agreement — was written off in 2022 and is now economical. *[INFERENCE]* This is a genuine, documented, recent change and it is the only one. *[FACT — CHAPTER 62]*

The coexistence model is proven in this market. AccessFintech's Synergy network — more than 250 institutions, more than a billion transactions a month — sits alongside source systems, integrates with Finastra's Fusion LenderComm, and has produced bank-reported outcomes of thirty per cent fewer fails and seventy-six per cent less email. *[FACT]* The architecture works here.

The structural argument from Chapter 41 is strong. A product sold to one institution, delivering value to that institution, escapes the standardization-failure mechanism entirely: no coordination, no adoption threshold, no cost-benefit misalignment across parties. *[INFERENCE]* This is the single best structural argument for the hypothesis, and it is why coexistence is the only class of approach with a clean theoretical path.

The incumbent's base is structurally receptive to exactly this pitch and no other. The moat is switching costs, not product love (Chapter 30). Switching costs protect against replacement and do nothing against addition. Dissatisfaction is real, back-loaded, and peaks at renewal in the analogous market. *[FACT]* Nobody has ever offered this base a low-risk optional exit. *[INFERENCE]*

And the channel exists. The implementation boutiques and systems integrators bill hours against the incumbent, already resell third-party technology into their installed base, and would find a coexistence layer expands their billable footprint rather than threatening it. *[FACT/INFERENCE]*

That is a serious case. It should be taken seriously before it is attacked.

The case against, in its strongest form

Now the attack. Each objection is stated as an opponent would state it, at full strength.

77.1 The coexistence position is already occupied

This is the strongest single objection and it is usually underweighted.

AccessFintech is a bank-backed, decade-old network doing coexistence in this market. More than 250 institutions. More than a billion transactions a month. A cap table containing JPMorgan, Citi, Goldman Sachs, BNY Mellon, Bank of America, Deutsche Bank and BNP Paribas. An integration with the incumbent servicing vendor's lender-data product since 2021, covering the syndicated-loan agent-to-lender data flow. A dedicated loan data management service. A private-credit partnership with a corporate trustee. And a November 2025 partnership connecting a major buy-side portfolio platform's global community to the network. *[FACT]*

If the entrant's core proposition is "sit alongside the incumbent core and reconcile loan data across counterparties," that is squarely this network's private-markets business, and it has the bank network, the incumbent integration, and the buy-side distribution. *[INFERENCE]*

Head-to-head on generic loan reconciliation is an uphill network-effects fight against an entity whose owners are the entrant's prospects. *[INFERENCE]*

The counter — that the incumbent network reconciles data but does not read contracts, does not govern provenance and never offers an exit — is real but thin. Reading contracts is table stakes (Chapter 64). Provenance is already being marketed by a data incumbent as source-cell traceability (Chapter 64). And the exit, per Chapter 77.4, may not be purchasable.

77.2 The access gate may simply not open

H3 is the hypothesis's most fragile claim and the one least supported by evidence.

The bank inherits liability for a third party's interception of non-public personal information and wire instructions. *[INFERENCE FROM THE APPLICABLE FRAMEWORK]* The regulatory trajectory across two jurisdictions runs from interception toward sanctioned, consented, revocable interfaces. *[FACT]* The academic literature says network-only inference is inferior to execution-informed inference and that encryption largely defeats it. *[FACT]* And the incumbent vendor can degrade an observer at will by changing its formats, with every incentive to do so. *[INFERENCE]*

There is no documented case of a bank granting a third-party vendor observation rights over its loan servicing traffic. *[UNKNOWN]* The hypothesis's central mechanism has no observed precedent in this market.

77.3 Messages are not state, and the reconstruction may be incomplete in ways that are invisible

H2 assumes a shadow can be built. Event interception captures deltas, and any state the servicing platform computes internally and never emits is invisible. *[FACT]* Historical positions predating the listener are invisible. Configuration-driven business rules are invisible. Delivery is at-least-once, so duplicates and losses occur. *[FACT]*

The failure mode is the dangerous one: a shadow that is ninety-eight per cent right and does not know which two per cent is wrong. *[INFERENCE]* Combined with the evaluation problem of Chapter 63 — no automated metric predicts extracted-specification quality — the institution is asked to trust a reconstruction that cannot be cheaply validated. Control functions do not buy that. *[INFERENCE]*

77.4 Progressive takeover requires a right the plan does not have

H5 is where the elegant story breaks.

The published guidance for the pattern states that it is unsuitable when requests cannot be intercepted and when the legacy system's source code cannot be modified — because disabling migrated features and redirecting internal calls requires modifying the legacy system. *[VENDOR CLAIM — THE PATTERN'S OWN DOCUMENTATION]*

A passive listener can shadow forever. Taking over a function requires a **write-path control point** that can suppress the incumbent and assert authority. That is not an extension of observation; it is a different contract, a different liability position, a different security review, and a different regulatory posture. *[INFERENCE]*

And Chapter 72 established that verification produces information, not authority. **There is no documented mechanism by which a shadow system acquires the contractual authority the register carries.** *[UNKNOWN]* The hypothesis's final step — "replacement becomes optional" — assumes a transition whose mechanism is unspecified.

77.5 Thirty years of evidence says the market adopts only what requires no change

Chapter 48's finding: the market will adopt anything that does not require it to change and will not adopt anything that does.

A shadow system that observes and reports requires no change and will be adopted — as a *reporting tool*, at reporting-tool prices. A system that takes over functions requires substantial change and will encounter Layers 2 through 5 of Chapter 30's moat: regulatory books and records, integration surface, operational muscle memory, and career risk. *[INFERENCE]*

The hypothesis may be self-defeating: the version that gets adopted is the version that is not valuable enough, and the version that is valuable enough does not get adopted. *[INFERENCE]* This is the sharpest form of the objection and the hypothesis has no clean answer to it.

77.6 The incumbent already sells the translation layer

The vendor ships interface products in which, per its own materials, it manages the translation and integration to the core. *[VENDOR CLAIM]* It can price, bundle, and integrate against a third-party translation layer, and it has both the relationship and the source-code access the entrant lacks. *[INFERENCE]*

77.7 The greenfield is being enclosed now

H6 is failing in real time, and Chapter 78 develops it.

77.8 Technology is at most half the problem

The legacy-displacement literature's headline finding: **"technology is at most only 50% of the legacy problem"** — organisational and process change is the other half. *[FACT]* A bank's operations, controls, audit and reconciliation teams must trust the shadow before any cutover, and that trust is earned in quarters, not sprints. *[INFERENCE]* It is a go-to-market cost, not an engineering one, and it is the cost most often omitted from plans of this type.

77.9 The emotional premise is not evidenced

Any version of the pitch that leads with customer resentment of the incumbent is leading with the least-supported claim in the entire evidence base. Three public reviews at 3.3 stars. *[FACT]* No documented pricing grievance. *[UNKNOWN]* No bank has ever publicly left. *[FACT]* The "would love to leave" proposition is an inference from a satisfaction-switching gap in an *analogous* market. *[INFERENCE]*

The clock

One development matters more than any other for the hypothesis's timing, and it deserves its own chapter.

On 30 April 2026 Versana closed a \$43 million round led by BNP Paribas, adding Apollo and Fitch Ventures as new investors alongside follow-ons from every existing bank owner, taking total raised above \$125 million — with an explicit buy-side and private-credit push. [FACT]

Read the components.

A mega-manager is now inside the consortium. The private-credit segment's largest participants are no longer merely prospective customers of a bank utility; one is an owner. [FACT]

The buy side is being monetised. The first buy-side subscriber came in January 2025, buy-side advisory committees have been formed, and a reconciliation module launched in January 2025. [FACT] The utility is moving from being an agent-data pipe to being a buy-side reconciliation product — which is the entrant's proposed territory.

Coverage is material. Reported active facility coverage exceeding \$4.1 trillion of notional. [VENDOR CLAIM]

Set alongside three further enclosure signals. The incumbent servicing vendor is keeping Loan IQ and explicitly aiming it at private credit. [FACT] A buy-side and CLO platform is in a sale process at up to \$3 billion, and a strategic acquirer could accelerate consolidation into private-credit operations. [FACT] A data and research incumbent already owns an AI-native private-credit monitoring platform positioned on provenance. [FACT]

The conclusion is uncomfortable and should be stated plainly. The "greenfield" framing of private credit was accurate in 2024 and is decreasingly accurate in 2026. The consortium model — which Part Seven

showed is the market's most effective institutional response and which compensates cost-bearers with ownership — is being extended into private credit deliberately, by the same actors, with a manager now inside the tent. *[INFERENCE]*

If this reading is right, H6 is the hypothesis's binding constraint, and speed to a defensible position is not a preference but the whole of the strategy.

[INFERENCE] If it is wrong — if private credit's fragmentation resists consortium formation because there is no agent chokepoint to pool (Chapter 60.2) — then the window is wider than it appears. This paper regards the first reading as more likely and flags the second as the most consequential place a reader might reasonably disagree. *[INFERENCE]*

78.1 How fast the clock actually runs

A clock chapter that does not say how to read the clock is a mood, not an analysis. Three correctives keep this one honest. *[INFERENCE]*

The base rate is slower than the announcements. The bank consortium's coverage was built agent by agent over roughly four years from a December 2022 go-live, gated by onboarding one institution at a time, and its own stated ambition — around ninety per cent of broadly syndicated loans — remains partial (Chapter 46). *[FACT / VENDOR CLAIM]* A funding round and a stated intention are not coverage. The private-credit push announced in April 2026 starts from the same constraint it took the syndicated business four years to work through, and in a segment with no agent chokepoint to onboard. *[INFERENCE]*

Announcement is not adoption, and this paper cannot see adoption. What is public is capital raised, owners added and modules launched. What is not public is which managers are live, on what, and whether the reconciliation module is used or merely available. That is test T6 of Chapter 79 and it is listed as an open question in Appendix E for the good reason that the answer determines whether this chapter is describing a closing window or a stated ambition. *[UNKNOWN]*

The observable indicators cut both ways. A reader who wants to know which reading is right should watch four things rather than the press: named private-credit managers going live rather than being announced; whether the reconciliation module is sold to the buy side as a product with its own price; whether a second mega-manager takes an ownership position; and whether the incumbent servicing vendor's private-credit push produces a reference customer. Enclosure completing looks like all four. Enclosure stalling looks like continued capital and continued silence on adoption. *[INFERENCE]*

The honest position is that the direction is evidenced and the rate is not. An entrant should plan for the fast case and would be wrong to assume it. *[INFERENCE]*

What would have to be true

The hypothesis survives only if a specific set of propositions holds. Each is stated so it can be tested within ninety days by someone with access to the market, and each has a named falsifier.

T1 — A named institution will grant observation access. *Test:* obtain a written statement of what a mid-sized private-credit manager's or agent bank's security and legal function would require to permit a third party to receive its loan servicing data flow. *Falsifier:* the requirement set includes conditions no vendor can meet, or the answer is a categorical no. This is the highest-risk proposition and should be tested first. *[HYPOTHESIS]*

T2 — Extraction accuracy on real agreements is high enough to be worth verifying. *Test:* run extraction against twenty real credit agreements with known ground truth, measured on the terms that drive money — margin grids, prepayment triggers, PIK mechanics, waterfall priorities — not on document-level metrics. *Falsifier:* accuracy on economically operative terms is low enough that human review costs more than manual encoding. *[HYPOTHESIS]*

T3 — A shadow can be built from available data. *Test:* attempt to reconstruct positions and accruals for one real portfolio from notices alone, and measure completeness against the institution's own records. *Falsifier:* the reconstruction is materially incomplete in ways not detectable from the notice stream — the Chapter 77.3 risk. *[HYPOTHESIS]*

T4 — Somebody will pay, and the buyer has a budget. *Test:* identify the actual purchaser. Efficiency buyers (operations) have small discretionary budgets; control buyers (risk, audit, compliance) have less discretionary but more defensible ones; loss-prevention buyers (credit, collateral) buy after an incident. *Falsifier:* every prospect agrees the problem is real and none has a budget line. This is the most common failure mode for products of this class. *[HYPOTHESIS]*

T5 — The value is realisable without takeover. *Test:* price the observe-and-verify capability alone, with no takeover path, and see whether it supports a business. *Falsifier:* the standalone capability prices as a reporting tool. If T5 fails, the hypothesis reduces to Chapter 77.5. *[HYPOTHESIS]*

T6 — The consortium has not already closed the segment. *Test:* determine, from the market rather than from press releases, what the bank-owned platform's private-credit product actually does today and which managers are live on it. *Falsifier:* the platform is already delivering reconciliation to the target buyer. *[HYPOTHESIS]*

T7 — Collateral verification is a separable, stronger proposition. *Test:* assess whether verifying that an asset exists and has been pledged once — the First Brands failure mode — can be delivered without observation access to a servicing core, using borrower-side and registry data. *Falsifier:* it cannot be done without the same access T1 requires. **If T7 holds and T1 fails, the collateral proposition survives the failure of the coexistence proposition, and this is the most important asymmetry in the whole test set.** *[HYPOTHESIS]*

Paths the hypothesis does not consider

The hypothesis is one of several structurally available approaches, and it is not obviously the best. Six alternatives deserve statement because the hypothesis is usually evaluated against nothing.

A1 — Sell the capability to the incumbents. The servicing vendor needs to modernise a large legacy core with dispersing expertise (Chapter 34) and is capital-constrained at sixteen-times leverage (Chapter 33). The administrators and outsourcers need better tooling to defend margins. The consultancies need something to sell alongside the platform. Each is a buyer of a capability rather than a competitor for a market. *Advantage:* no access gate, no anchor problem, immediate distribution. *Disadvantage:* a supplier's economics, not a platform's, and dependency on a single buyer — which is the Symbiont risk (Chapter 52). *[INFERENCE]*

A2 — Become the administrator. Private credit's operating problem is being solved today by third-party administrators charging for labour. A technology-enabled administrator sells an outcome rather than a tool, has legitimate access to the data because it is performing the function, and faces no interception gate at all. *Advantage:* the access problem dissolves; the buyer has an existing budget line. *Disadvantage:* a services business with services economics, competing against entrenched administrators with scale. *[INFERENCE]*

A3 — Verify collateral only. The narrowest and, on the evidence of Chapter 69 and Chapter 73, the best-evidenced proposition: no party is authoritative on whether an asset has been pledged into more than one lender's facility; the 2025–26 frauds reached named institutions through exactly that gap; and one entrant is already working the adjacent problem in asset-backed lending. *Advantage:* a recent, attributable failure at named institutions; no contest with an existing authority; a loss-prevention buyer. *Disadvantage:* a

narrower market than the headlines imply — Chapter 73 shows that most of what failed was diligence the institutions could have performed themselves — and the same access question in a different form. *[INFERENCE]*

A4 — Build for the control function, not the operations function. Sell independent recomputation as a control artefact to risk, audit and compliance rather than as an efficiency tool to operations. *Advantage:* a supervisory tailwind (Chapter 56) and a buyer whose budget is defensible. *Disadvantage:* a smaller contract, and the risk of being classified as an audit tool rather than infrastructure. *[INFERENCE]*

A5 — Go through the channel. Sell to and through the implementation boutiques and systems integrators who already sit inside these banks (Chapter 13). *Advantage:* distribution into a trusting relationship, and an incentive alignment that expands rather than threatens their billing. *Disadvantage:* the channel captures the margin and the customer relationship. *[INFERENCE]*

A6 — Decline. The strongest arguments for declining are in Chapter 77 and are not weak. *[INFERENCE]*

Two observations about this list.

First, most of these alternatives dissolve rather than solve the access gate of Chapter 74, which the coexistence hypothesis must confront directly. A2 in particular obtains data access by performing the function that generates the data — which is the only clean solution to the problem the hypothesis treats as its central obstacle. *[INFERENCE]*

Second, they are not mutually exclusive, and A3 combined with A4 is a coherent narrow strategy that requires none of the hypothesis's riskiest claims: it needs neither H3 (observation access to a servicing core) nor H5 (takeover), and it inherits the strongest single piece of recent evidence in the paper. *[INFERENCE]*

Verdict on the central question

The paper's organizing question was what has changed between the last generation of modernization attempts and 2026 that could make a structurally different outcome possible.

81.1 What has genuinely changed

Four things, and only four. *[INFERENCE]*

One: interpretation became economical. Reading bespoke credit agreements and unstructured notices at scale was not economically possible in 2022 and is now. This is documented in peer-reviewed research and in production tooling, and it is the only change in technical capability that matters. *[FACT / INFERENCE]*

Two: governance acquired a regulator. Between 2024 and 2026 the FSB, the Bank of England and the SEC all named valuation opacity and data gaps in writing. *[FACT]* This converts a governance capability from an efficiency purchase to a compliance purchase.

Three: the failure mode changed, and the new one is addressable. The market's canonical operational failures were payment errors — Revlon, the \$81 trillion near-miss. Its 2025–26 failures were **collateral and data-integrity frauds**: First Brands, with more than \$9 billion of liabilities, and Tricolor, at \$900 million-plus, both driven by double-pledged collateral and fabricated receivables, with losses and exposures reported at several major institutions. *[FACT]* That is a categorically different problem, it maps to a row in the authority table where **no authority exists** (Chapter 69), and it is recent, attributable and expensive.

Four: the human layer began to disappear. One-to-two-year training against twenty-to-seventy-per-cent attrition *[VENDOR CLAIM]*; 2026 AI-attributed back-office layoffs at scale *[FACT]*; and the platform expertise

dispersing to consultancies (Chapter 34). The manual model is losing the people who make it work. *[INFERENCE]*

81.2 What has not changed

Five things, and they are the ones that killed every previous attempt. *[INFERENCE]*

The agent still holds the register and is still contractually the truth. *[FACT]*
The agreement is still bespoke in every economically significant term.
[FACT] Settlement still requires consent from parties not compensated for speed. *[FACT]* Transparency still transfers value from the parties who control the infrastructure. *[INFERENCE]* And the banks still respond to any valuable infrastructure layer by consortium-ising it — now with a mega-manager inside the consortium. *[FACT]*

In the Chapter 26 framework: **Category 2 changed. Categories 3 and 5 did not. And the market's failures have always been in Categories 3 and 5.** *[INFERENCE]*

81.3 The verdict

On the strong hypothesis — observe, shadow, take over, make replacement optional — the public evidence does not support it. The precise form of that finding matters. *[INFERENCE]*

H3 and H5 are unsupported rather than falsified. No documented case exists of an institution granting a third party observation access to its loan servicing traffic, and the regulatory trajectory across two jurisdictions runs from interception toward sanctioned interfaces. *[FACT / UNKNOWN]* That is an absence of precedent plus an adverse trend — not a demonstration that access cannot be granted. An entrant that obtains it has not defied a law of nature; it has done something for which this paper could find no prior instance. *[INFERENCE]* Likewise the takeover step requires a write-path control point and, ultimately, contractual authority that verification

does not confer, and **no mechanism for acquiring that authority has been specified by anyone, including this paper.** *[UNKNOWN]* Unspecified is not impossible.

The self-defeat problem is the more serious objection, because it does not depend on an absence of evidence: the version of the product that clears the access and change-management gates is the observe-and-report version, and observe-and-report prices as a reporting tool (Chapter 77.5). *[INFERENCE]*

On a narrower hypothesis — verification and governance of the things about which no party is authoritative — the evidence is materially stronger. But the two candidate problems must be separated, because they are not the same business. *[INFERENCE]*

Collateral integrity. There is a verifiable external fact: an asset exists or does not, and has been pledged once or more than once. The 2025–26 frauds reached named institutions through that gap, in documented amounts.

[FACT] The buyer is a lender, and the purchase is loss prevention, which is the least discretionary spending category in a credit institution after regulatory compliance. *[INFERENCE]* It contests no one's authority, requires no observation access to a servicing core, and needs no takeover path.

The open question that stood here in an earlier draft — how much of this ground existing security-interest filing regimes already cover — is now answered, and the answer narrows the finding in two ways that a reader should carry forward. **The filing system covers rank, not existence, and was designed that way; and in every case examined an authoritative verification mechanism existed and was not used.** *[FACT / INFERENCE]* The residual gap that no single institution can close alone is therefore narrower than the failures make it look: it is the question of whether an asset has been pledged into another lender's facility, and nothing else. Chapter 73 and Appendix F set out the mechanics and the qualifications.

Private-credit valuation. There is no verifiable external fact — in the absence of a market price, a valuation is a considered judgment — so "verification" here means auditing a process rather than establishing an answer. [INFERENCE] The supervisory tailwind is real and named. [FACT] But the natural buyer is the manager whose judgment is being examined, which is a structurally conflicted purchase, and the parties with the strongest interest in the answer — limited partners and supervisors — do not buy software. [INFERENCE] This is a genuine problem with an awkward buyer, and the paper's earlier draft erred in grouping it with collateral as a single opportunity.

Both share the property that they escape the standardization-failure mechanism of Chapter 41, because each is sold to one institution for that institution's benefit. [INFERENCE] Only the first has a buyer whose incentives point the same way as the product's.

81.4 The residual uncertainty, stated honestly

Three things could overturn this verdict, and a reader who believes any of them should reach a different conclusion.

If observation access turns out to be routinely obtainable — if some institution has already granted it and the paper simply could not find the case — then H3 survives and the strong hypothesis is materially better than assessed here. [UNKNOWN] This is the single most valuable thing a reader with market access could establish, and the paper's judgment on Path A should be read as provisional on it.

If private credit resists consortium formation because there is no agent chokepoint to pool (Chapter 60.2), then the window is wider than Chapter 78 assumes and speed is less critical. [HYPOTHESIS]

If a downturn deepens materially, the urgency effect may dominate the budget-freeze effect, and institutions may buy governance capability they would otherwise defer. [HYPOTHESIS] Which effect dominates is not predictable from outside. [UNKNOWN]

81.5 Where the verdict sits against the hypothesis as stated

The hypothesis was decomposed in Chapter 75.4 precisely so that it could fail in parts. It did. Recording which parts, and on what evidence, is what makes the verdict a finding rather than an impression.

CLAIM	STATUS ON THE PUBLIC EVIDENCE	WHERE
H1 Interpretation is good enough to be worth verifying	Supported , with the evaluation problem attached: no automated metric predicts extracted-specification quality, so human review is a permanent cost line rather than a transitional one	Ch. 62, 63; T2
H2 The shadow can be reconciled continuously enough to be trusted by control functions	Unresolved . Plausible in architecture, unproven on completeness, and the incompleteness would be invisible from the notice stream	Ch. 77.3; T3
H3 An institution will grant observation access	Unsupported — and this is an absence of precedent, not a demonstration of impossibility . No documented case exists; the regulatory trajectory in two jurisdictions runs away from interception	Ch. 74; T1
H4 The capability is worth enough to one institution to be bought	Contested . The buyer with the budget wants a control artefact; the buyer who wants the efficiency has the smallest discretionary budget in the institution	Ch. 77.5; T4
H5 Progressive takeover is achievable	Unsupported, and no mechanism has been specified by anyone, including this paper . Takeover requires a write-path control point and, ultimately, contractual authority that verification does not confer	Ch. 74.4, 81.3; T5
H6 The segment remains unclaimed long enough	Direction evidenced, rate unknown . Enclosure is deliberate and under way; adoption is not publicly observable	Ch. 78, 78.1; T6

Two readings of that table are available and only one of them is supported.

The unsupported reading is that H3 and H5 are refuted. They are not. They are propositions for which this paper could find no instance, in a market that does not publish its arrangements. An entrant who obtains observation access has not defied a structural law; it has done something undocumented. *[INFERENCE]*

The supported reading is narrower and harder to escape: **the hypothesis requires all six claims, and two of the six rest on an absence of evidence while a third — self-defeat — does not**. *[INFERENCE]* A conjunction of six is

not made likely by the fact that each is individually arguable, and the one that does not depend on missing evidence is the one that bites.

And the asymmetry recorded at T7 is what carries into the next parts: **the collateral proposition survives the failure of H3 and H5**, because it requires neither observation access to a servicing core nor a takeover path. That is why Part Thirteen assesses it separately rather than as a component of the coexistence path.

81.6 The honest summary

Something has changed, it is real, and it is smaller than the enthusiasm suggests.

AI made it possible to read this market's documents. That is a genuine unlock of a capability the market had written off, and it is the first such unlock in a decade. It did not make it possible to obtain the agent's consent, to acquire contractual authority, to change what a bank's audit function will accept, or to prevent seven institutions that own the market's infrastructure from building the same thing and owning it themselves.

The opportunity that survives contact with all of the evidence is narrower than the one usually described, sits in a different part of the market, and is defended by different things.

Two parts follow, and they answer different questions. Part Twelve asks what **the industry** could change, given everything Parts One through Ten established about why it has not. Part Thirteen asks what a **new entrant** could profitably do. The answers do not coincide, and the gap between them is itself one of this paper's findings: a market can have a rational direction that no single company can profitably push it toward.

XII

PART TWELVE

What the industry could actually change

Ten parts of description and one part of hypothesis-testing have earned the right to a prescription. This part supplies one — for the industry, not for a vendor.

It is written for the readers the front matter named and the previous draft of this paper did not serve: the trade-association executive, the global head of syndicated lending, the agent-bank operations head, the CLO manager, the private-credit chief operating officer, the supervisor, and the lawyer who will be asked whether any of it is possible. Each of them holds a piece of this problem. None of them can sell software to the others, and none of them needs to.

Every recommendation below is derived from Parts One through Ten. Nothing new is invented, no solution is proposed that the body of the paper has not already shown to be structurally available, and where the analysis implies that a thing cannot be done, this part says so rather than recommending it anyway.

The distinction this part rests on

There are two different questions and this paper has been asked both.

"What should this market do?" is a question about collective rationality. It asks what arrangement of documents, standards, incentives, authority and technology would leave the participants collectively better off, and it is answered by reasoning about mechanisms.

"What business could someone build here?" is a question about appropriability. It asks whether any single firm can capture enough of the value of an improvement to fund making it, against competitors who own the customers.

These come apart in this market more sharply than in most, and the reason is Chapter 41. The market's characteristic failure is that costs and benefits fall on different parties. That fact makes some improvements collectively rational and individually unfundable: the improvement is worth more than it costs, and nobody who could make it captures enough to justify making it.

So Part Twelve says what the industry could do. Part Thirteen asks separately whether anyone can profitably make it happen, and reaches a narrower answer. A reader who finds the two parts in tension has understood them correctly. *[INFERENCE]*

One further framing point. The recommendations below are ordered by how much coordination they require, from none to industry-wide, because coordination cost is the binding constraint identified throughout the paper. **Chapter 91's list — what a single institution can do alone — is where a reader with limited patience should start, because it is the only category the paper's own analysis says will actually happen.** *[INFERENCE]*

What should remain bespoke

Prescriptions about this market almost always begin with what should be standardized. Beginning there is an error, because it invites the standardization of things whose variability is doing real work.

Chapter 36 established the boundary and it should be treated as fixed. A credit agreement's commercially significant terms — mandatory prepayments, business undertakings, financial covenants, representations and warranties, conditions to drawdown — are the negotiated allocation of risk between a borrower with leverage and lenders with capital. *[FACT]* A covenant package is a price. A margin grid is a price. A prepayment provision is a price. Standardizing them means standardizing the outcome of a negotiation. *[INFERENCE]*

The following should therefore remain variable, and any initiative that proposes otherwise should be resisted by the parties whose economics depend on the variation:

Credit economics. Pricing, margin grids, fee menus, original issue discount, ticking-fee structures. These are how risk is priced and how competition between lenders expresses itself.

Covenant architecture. Maintenance versus incurrence, levels, cushions, baskets, builder provisions, and the entire apparatus of negotiated flexibility. The erosion of private credit's covenant advantage (Chapter 55) is a competitive outcome, not a standardization failure.

Risk allocation. Prepayment mechanics, excess-cash-flow sweeps, mandatory prepayment triggers, security packages, guarantor structures.

Transfer restrictions — with one qualification. Consent regimes and disqualified-lender lists are the direct cause of the market's settlement times (Chapter 22.3), and the trade association's own diagnosis is that "changing credit agreement provisions is necessary to shorten settlement

times." *[FACT]* It is nonetheless the borrower who wants control over who holds its debt, and that is a legitimate commercial interest rather than an operational defect. The realistic prescription is not to abolish transfer restrictions but to **standardize their machinery while leaving their substance negotiated**: a common format for a disqualified-lender list, a common deemed-consent period, a common electronic channel for requesting and recording consent. Those are exchange and workflow standards, not economic ones. *[INFERENCE]*

A note on jurisdictional divergence. New York-law and English-law documentation have not converged in three decades and the Asia-Pacific framework has diverged further (Chapter 36.3). *[FACT]* This paper does not recommend convergence, because it is not achievable and because the divergence reflects genuinely different legal systems. Chapter 85 proposes the alternative.

The general principle: standardize the container, negotiate the contents.

[INFERENCE] Almost every successful standard in this market's history conformed to it, and almost every failure violated it.

What could usefully be standardized, and at which layer

The word "standard" hides at least six different things in this market, and the paper's evidence says they have completely different prospects. Distinguishing them is the most useful contribution this part can make, because the market's standardization debates routinely conflate a layer that has already failed with one that has never been attempted.

84.1 The six layers

Layer 1 — Identifier standards. A common name for a loan, a facility, a tranche, an entity. **Status: attempted, failed, and correctly abandoned in favour of translation.** Four identifier systems coexist; two competing entity identifiers exist; the trade association's answer was a cross-reference service. *[FACT]* Chapter 37.3 explains why: identifier fragmentation is a competitive equilibrium among vendors who each own a scheme, not a coordination failure among users. *[INFERENCE]* **Recommendation: stop trying. Invest in translation (Chapter 85).**

Layer 2 — Exchange standards. A common envelope for transmitting a message: the schema, the transport, the acknowledgment. **Status: attempted, published, not adopted.** *[FACT]* Chapter 38 documents the failure and its four causes. **Recommendation: exchange standards are worth having but will not be adopted through voluntary mutual agreement, because the cost falls on the agent and the benefit on the lenders. They become adoptable only through the mechanisms in Chapter 86.** *[INFERENCE]*

Layer 3 — Semantic standards. A common *meaning*: an agreed definition of "facility fee," of what fields constitute an interest period, of what a rate-set notice must assert. **Status: never seriously attempted, and the actual problem.** The market's own literature identifies it precisely — one bank sends start date and end date where another sends start date and tenor; "facility fee" means commitment fee to one sender and utilisation fee to

another. *[FACT]* **This is the highest-value standardization work available to this market that has not been seriously attempted, and its adoption economics are materially better than those of the identifier and exchange-standard efforts examined in this paper.** *[INFERENCE]*

That last point deserves precision, because the temptation is to overstate it.

A shared definitional register — what the market means by each term, maintained by the trade bodies, versioned, and referenced rather than embedded — does not inherently require any of the things that killed the other efforts. It does not require replacing a servicing system. It does not require a vendor to surrender an identifier franchise. It does not require changes to negotiated credit economics. It does not transfer contractual authority. It does not need universal adoption to be useful, since a register has value to the first party that consults it. And it does not require one institution to fund a major operational change principally for another institution's benefit. *[INFERENCE]*

It would nonetheless require real work from real parties: governance and a maintaining body, versioning, mapping from each institution's existing usage, implementation in the systems that consume the definitions, regional and legal interpretation where New York-law and English-law usage diverge, and change management. None of that is free.

So the correct statement is that the Chapter 41 cost-benefit asymmetry is materially weaker at the semantic layer, not that it is absent. *[INFERENCE]*

That is a smaller claim than the one this chapter's first draft made, and it is the one the evidence supports. It is also still enough to distinguish this layer from every other: the asymmetry that defeated the messaging standard was structural and unavoidable, while the costs above are ordinary programme costs of a kind trade bodies already absorb for documentation work.

Layer 4 — Authority standards. An agreed statement of who is authoritative about what. **Status: implicit in every credit agreement, never articulated as a standard.** Chapter 69's table is an attempt to write down what the market already practises. *[INFERENCE]* Making it explicit — a published map of which party is authoritative for which data element, and what a downstream party's recourse is when it disagrees — costs nothing, changes no system, and would prevent a great deal of wasted engineering aimed at building golden sources that the market's contracts do not permit to exist. **Recommendation: worth doing, cheap, and nobody is doing it.** *[INFERENCE]*

Layer 5 — Evidence and provenance standards. An agreed format for *how a party knows what it asserts*: the clause a term came from, the version of the agreement, whether a human or a model read it, who reviewed it, when it changed. **Status: emerging commercially, unstandardized.** One data incumbent already markets "source-cell traceability" in private credit. *[FACT]* Chapter 71.1 sets out what a provenance record must contain.

Recommendation: this is the layer where a standard would create the most durable value, and where commercial adoption is running ahead of any standard — which historically is the moment at which a trade body can still shape it and shortly after which it cannot. *[INFERENCE]*

Layer 6 — Workflow standards. Agreed sequences and deadlines: how long a consent request may sit, what constitutes a complete assignment package, when a deemed consent operates. **Status: partially attempted through settlement conventions, with results measured in single-digit percentage improvements over twenty-five years.** *[FACT]* **Recommendation: this is where the trade bodies are already concentrating, and Chapter 39.1 explains why progress is slow — the bottleneck lives in negotiated contract provisions rather than in conventions.**

84.2 The ranking

Ordered by expected value net of coordination cost, on this paper's evidence:

RANK	LAYER	WHY
1	Semantic	High value, never seriously attempted, and the layer where the cost-benefit asymmetry is weakest — not absent <i>[INFERENCE]</i>
2	Provenance / evidence	High value, commercially live, standardizable now and probably not later <i>[INFERENCE]</i>
3	Authority	Low cost, prevents category errors in every downstream project <i>[INFERENCE]</i>
4	Workflow	Real value, already the focus, structurally slow <i>[INFERENCE]</i>
5	Exchange	Real value, blocked by the incentive problem until Chapter 86 is solved <i>[INFERENCE]</i>
6	Identifier	Abandon; translate instead <i>[INFERENCE]</i>

The striking feature of that ranking is that the market's standardization effort has historically concentrated on layers 5 and 6 — the two with the worst prospects — and has never seriously attempted layers 1 through 3 of the ranking. *[INFERENCE]* That is not a criticism of the trade bodies, whose mandate came from a secondary-trading agenda in the 1990s where identifiers and messages were the obvious targets. It is an observation that the agenda has not been revisited since the problem changed.

What should be translated rather than unified

The identifier history contains a general principle that the market has discovered accidentally and has not generalized.

The principle: where multiple standards persist because their owners have competing commercial interests in them, translation between standards is achievable and unification is not. *[INFERENCE]*

The evidence is the Loan Identifier Cross-Reference. Faced with four identifier schemes whose owners would not abandon them, the trade association did not attempt a fifth. It built a mapping service. *[FACT]*

That is a concession, and it is also the correct engineering response to a competitive equilibrium.

85.1 Does it generalize?

Testing the principle against the six layers of Chapter 84:

Identifiers — yes, and already done. *[FACT]*

Jurisdictional documentation — yes, and partially done. The bodies produced a comparative users' guide and a shared glossary rather than a merged document. *[FACT]* The paper's recommendation is to extend this deliberately: a maintained concordance between New York-law, English-law and Asia-Pacific provisions, at the level of defined terms and operative mechanics, would let systems and lawyers move between frameworks without pretending they are the same. *[INFERENCE]*

Semantics — partially, and this is the interesting case. A definitional register (Chapter 84, layer 3) could be built either as a *unified* dictionary that everyone adopts, or as a *translation* layer that maps each sender's usage onto a reference meaning. The second is materially easier, because it asks nothing of the sender's systems, and it is what an interpretation layer built

on modern language models is naturally good at (Chapter 62). It still requires someone to build and maintain the reference meanings and the mappings, which is where the real cost sits. **Recommendation: build the reference meanings, and translate to them, rather than requiring anyone to adopt them.** *[INFERENCE]*

Exchange formats — yes, and this reframes the failed messaging project. The FpML effort failed partly because it required senders to emit a common format. *[FACT]* A translation approach requires senders to change nothing and places the normalization burden on the recipient or on a shared service. That is more expensive in aggregate and materially more achievable, because it moves the cost onto the party that captures the benefit — which is the asymmetry Chapter 41 identifies as fatal, corrected rather than eliminated. *[INFERENCE]*

Authority — no. Two parties cannot both be authoritative about the same fact and be translated between. Authority must be allocated, not mapped. *[INFERENCE]*

Provenance — no, and this is the one place a genuine standard is still available. A provenance record is new; there is no incumbent format to compete with, and therefore no equilibrium to be trapped by. *[INFERENCE]*

85.2 The general rule

Translate where incumbency has already created competing standards. Standardize where nothing exists yet. Allocate where the question is authority rather than format. *[INFERENCE]*

This is the single most transferable principle in Part Twelve, and it inverts the instinct of most standard-setting bodies, which is to attempt unification first and settle for translation only after failure. In this market the sequencing should be reversed on the layers where incumbents already exist.

The agent-incentive problem, and the five ways to solve it

Chapter 41 identified the mechanism that has defeated this market's standardization efforts: **the costs fall on the agent and the benefits fall on the lenders**. This chapter asks what could realistically change that.

Five classes of mechanism exist. The paper's evidence permits an assessment of each.

86.1 Fees

The mechanism. Pay the agent for the service the market wants: a per-notice fee for structured emission, an increased agency fee for a defined data-service standard, a fee for same-day consent processing.

Assessment: available, underused, and the most obvious unexplored option. *[INFERENCE]* The market already prices agent services — the \$3,500 to \$4,500 assignment fee is set unilaterally by the agent and waivable at its discretion. *[FACT]* The machinery for paying the agent exists; what does not exist is any fee tied to service *quality* rather than to transaction occurrence. **A settlement-speed or data-quality component in the agency fee would align the cost-bearer with the benefit, which is precisely what Chapter 41 says is required.** *[INFERENCE]*

The objection, which is serious. The agency fee is negotiated between the borrower and the agent, while the beneficiaries of faster service are the lenders. *[FACT]* So a fee mechanism requires the *lenders* to pay the agent, which is a change to the market's fee architecture rather than a change to a fee level. *[INFERENCE]* The delayed-compensation regime is the closest existing precedent for lenders compensating each other for timing, and it explicitly exempts agent-caused consent delay. *[FACT]* Extending something like it to cover agent processing time is the concrete version of this recommendation.

86.2 Ownership

The mechanism. Give the agent equity in the utility that consumes its data.

Assessment: proven, and already the market's most successful

modernization. *[INFERENCE]* This is exactly what Versana does: the agent banks are founders and shareholders, so the value the platform creates accrues partly to the party bearing the cost of supplying it. *[FACT]* Chapter 41's Corollary 2 predicted that initiatives which pay the agent can succeed, and this is the demonstrated case.

The limitation. Ownership works when the cost-bearers are few and large enough to form a consortium. It scales to the top ten agent banks and does not obviously scale to the long tail, to non-bank agents, or to private-credit managers acting as their own agents. *[INFERENCE]* And Part Seven's analysis applies: a consortium solves shared-infrastructure problems with symmetric benefits and cannot solve distributional ones.

86.3 Commercial incentive

The mechanism. Make good data a competitive advantage in winning agency mandates.

Assessment: theoretically the cleanest solution and currently absent.

[INFERENCE] If borrowers or lenders selected agents partly on operational service quality, the agent would invest without needing to be paid separately. The reason this does not happen is that **agency is generally not competed for on its merits — it follows the arranging mandate**, which is won on price, distribution and relationship. *[INFERENCE]*

What would have to change. Agency would have to be unbundled from arranging, or at least separately assessed. There is a live precedent: independent third-party agents already market themselves on the conflict-of-interest and service arguments (Chapter 9, Chapter 11), and the private-credit market uses non-bank agents routinely. *[FACT]* **The recommendation**

is modest and real: lenders who care about operational quality should say so at syndication, and should treat the identity and service standard of the agent as a term rather than an afterthought. *[INFERENCE]*

86.4 Contractual reform

The mechanism. Change what the credit agreement requires the agent to do and by when.

Assessment: the trade association's own choice, and correctly so. Its 2026 programme targets credit-agreement provisions — consent tracking, agreement updates, lender communications — precisely because it concluded the bottleneck lives there. *[FACT]* Deemed-consent provisions, standardized disqualified-lender machinery, and defined processing windows are contractual, not technological, and they are negotiated at origination when the borrower and arranger have the pen.

Assessment of prospects: slow but genuine. *[INFERENCE]* Contract terms propagate at the speed of refinancing, which in this market means a meaningful share of the stock turns over within a few years. This is the mechanism most likely to produce measurable change in settlement times, and it is the one with the longest lag.

86.5 Regulatory compulsion

Reserved for Chapter 87, because the evidence supports it in one place and not generally.

86.6 The recommendation

Ownership is proven and does not scale beyond the largest agents.
Contractual reform is slow and real. Fees are available and unexplored, and are the mechanism the paper's own analysis most directly implies.
Commercial incentive requires unbundling agency from arranging, which nobody is proposing. *[INFERENCE]*

An industry serious about this problem would attempt all four, and would recognise that the reason none of them has been attempted is not that they are difficult but that **no participant's mandate covers the gap between the cost-bearer and the beneficiary**. That is what a trade body is for, and it is a different agenda from the one the trade bodies currently run. *[INFERENCE]*

Where regulation is actually required

This paper does not default to regulation. Most of what is wrong with this market is a coordination and incentive problem that regulation would address clumsily, and the market's one large recent deregulation — the December 2025 rescission of the leveraged-lending guidance — was justified by the agencies on grounds this paper's Part One found credible. *[FACT]*

But the analysis identifies one problem that the paper's own reasoning says voluntary action cannot solve, and integrity requires naming it.

87.1 The case: post-trade transparency

The argument, assembled from Chapters 21, 49 and 53.

Loans are not securities and are not reported to any transaction-reporting facility; there is no consolidated tape. *[FACT]* Marks are quote-derived rather than trade-derived. *[FACT]* The corporate-bond evidence indicates that post-trade transparency substantially reduced the bid-ask spreads investors pay, though the transfer of that finding to loans is inference rather than demonstration (Chapter 53). *[FACT / INFERENCE]* Transparency would therefore move value from dealers to investors. The dealers own the market's infrastructure ventures. *[FACT]* A consortium of dealers has not built and will not build a facility that costs dealers money, and the trade association's 2026 reform programme contains no transparency mandate. *[FACT]*

Conclusion: this is the one problem in the paper for which the evidence supports the claim that no voluntary configuration produces the outcome.

[INFERENCE] If post-trade transparency in loans is desirable, it will require compulsion, as it did in corporate bonds.

The honest counter-argument, which a regulator should weigh. The academic literature also holds that transparency can make dealers reluctant to carry inventory. *[FACT]* Loans are less fungible, trade far less often, and have a smaller dealer population than corporate bonds, so the

liquidity cost could be larger and the spread benefit smaller. *[INFERENCE]*

This paper does not claim that a loan tape would be net beneficial. It claims that **if it is beneficial, it will not arrive voluntarily** — which is a different and better-supported statement. *[INFERENCE]*

87.2 The second candidate: private-credit valuation process

Weaker, and worth stating carefully.

Supervisors have named valuation opacity and data gaps in writing, across three institutions and two jurisdictions. *[FACT]* Private-credit fair value has no external authority: the manager produces the number. *[FACT]* Chapter 69 established that this is a **process** problem rather than a fact-finding problem — there is no external answer to check against — and Chapter 81.3 established that its natural buyer is the party being examined. *[INFERENCE]*

That combination — no external fact, and a conflicted buyer — is the classic signature of a problem requiring an external standard rather than a market solution. *[INFERENCE]* The instruments available are ordinary and already partly deployed: examination priorities, disclosure requirements, and reporting thresholds. *[FACT]*

The defensible regulatory object is the process and its evidence — methodology governance, inputs, consistency of application, independence of review, and the auditability of the chain from assumption to number — not the number itself. *[INFERENCE]*

Four things this paper explicitly does **not** claim, because each is a step the argument does not support and each would be a fair objection if left implicit:

It does not recommend that regulators prescribe a fair value. Valuation here is a judgment about future cash flows in the absence of a market. Mandating a methodology would produce uniform numbers without producing accurate ones, and would convert a judgment into a compliance artefact. *[INFERENCE]*

It does not claim that software can determine the objectively correct value of an illiquid loan. There is no such value to determine (Chapter 69). What a system can do is record how a valuation was reached, whether the method was applied consistently across the portfolio and across periods, and what changed between marks. That is evidence about a process, not an answer.

It does not treat valuation disagreement as error. Two competent managers can mark the same asset differently and both be defensible; a difference between them is not a break in the sense Chapter 69 uses the term. This distinction matters operationally, because a system that flags every valuation divergence as an exception produces noise.

It does not assume that supervisory attention creates a buyer or a budget. Whether examination priorities convert into procurement is a separate empirical question this paper cannot answer, and Chapter 94.2 states plainly that the natural buyer here is structurally conflicted. *[UNKNOWN]*

The recommendation is therefore deliberately weaker than the transparency case in §87.1. That one rests on a demonstrated structural conflict and a documented absence of any voluntary path. This one rests on an inference from the combination of no external fact and a conflicted buyer, and it should be read as a direction supervisors may reasonably take rather than as a demonstrated necessity. *[INFERENCE]*

87.3 Where regulation is not required

Explicitly, on this paper's evidence: **identifiers** (a competitive equilibrium best addressed by translation), **documentation** (the variation is economically purposeful), **settlement speed** (the bottleneck is negotiated contract provisions the parties can change themselves), and **messaging formats** (an incentive problem with four non-regulatory solutions in Chapter 86). *[INFERENCE]*

What AI should and should not be allowed to do

Chapter 65 drew the boundary for a system designer. This chapter states it as an industry standard, because the boundary will be drawn either deliberately or by the first serious incident.

88.1 The rule

AI interprets. Deterministic systems calculate and transact. Governance controls the crossing between them. *[INFERENCE]*

The three parts are not negotiable independently. Interpretation without deterministic execution puts a probabilistic system in the payment path. Deterministic execution without interpretation is the 2022 status quo, in which the encoding burden pushes work onto spreadsheets. Either without governance is an unauditable financial calculation.

88.2 What belongs on each side

Deterministic, always: monetary calculations, accruals, balances, positions, the register, payment allocations, ledger entries, entitlements, transaction state. These are legally operative determinations that move money and create or extinguish obligations, and the system must be able to say not only what the number is but why — reproducibly, and in a form that survives audit and litigation. *[INFERENCE]*

Probabilistic, with verification: reading a bespoke agreement to propose a structured representation; classifying and extracting from an unstructured notice; resolving semantic ambiguity; explaining a discrepancy and proposing its cause; drafting; triaging exceptions; capturing expert knowledge. Each produces an *input to a decision* rather than the decision. *[INFERENCE]*

88.3 The crossing, which is the part that needs a standard

The interesting requirement is not the two lists. It is what must accompany a proposed interpretation for a deterministic system, an auditor and a supervisor to accept it: the clause it came from, the structured assertion, the method that produced it, the confidence, the human who reviewed it, the version of the agreement, the date, and the full history of changes. *[INFERENCE]*

That is a provenance record, and Chapter 84 identified it as the layer where a genuine standard is still available. *[INFERENCE]* An industry that agrees what a crossing record must contain — before every vendor invents its own — obtains auditability, comparability across institutions, and a defensible answer to a supervisor asking how a number was produced. An industry that does not will get one such format per vendor, and will be translating between them in ten years.

88.4 The evaluation obligation

One further industry-level point, from Chapter 63. **No automated metric currently predicts the quality of an AI-extracted specification, and human evaluation remains necessary.** *[FACT]* Any institution deploying interpretation at scale therefore takes on a *measurement* obligation, not merely an accuracy target: it must be able to state how it knows the extraction is right, on a sample, continuously, with the parallel run as the only demonstrated proving mechanism. *[INFERENCE]* Institutions that skip this will discover their error rate through their counterparties.

88.5 The rule restated for the hardest case

The collateral domain of Chapter 73 is where the boundary is most likely to be crossed by accident, because the value of machine reading there is real and the temptation to over-claim is correspondingly strong. Stated for that case: a model can read every invoice rather than a sample, resolve obligors across facilities, and identify the contradictions that should trigger a confirmation. It cannot confirm the obligation — only the account debtor

can — and it cannot see into another lender's facility at all. **AI can detect a reason to doubt; it cannot manufacture authoritative evidence that does not exist.** *[INFERENCE]* An institution that treats a model's flag as a finding has moved a Category 2 capability into a Category 5 role, which is the error this chapter exists to prevent. Appendix F §P sets out the pipeline and its limits.

FIGURE 15

From proposition to institutional action



A model's output is a proposition, not an institutional fact: it acquires standing only through provenance, challenge, verification and the authority to act. §88.3 sets the provenance record this requires. Supervisory statements shown are quoted from the sources cited in Part Nine. [FACT]

The authority model the industry should accept

This is the chapter that carries the industry agenda, and its recommendation is a refusal as much as a proposal.

89.1 Stop pursuing a single golden source

The market does not lack a golden source. **Authority is distributed by design, deliberately, by contract, and different questions have different correct answer-holders.** [INFERENCE] The agent is authoritative on the register; the borrower on its own condition; the custodian on where the cash is; the trustee on the waterfall; the administrator on net asset value; the manager, in private credit, on valuation; and nobody on market price.

[FACT / INFERENCE PER CHAPTER 69]

Three decades of evidence say the market will not converge on one source: identifiers did not converge, documentation did not converge, messaging did not converge, and the most successful centralization to date is partial, voluntary, gated by agent onboarding and compensated by ownership. [FACT / INFERENCE]

An industry that keeps funding golden-source projects is funding a category error. [INFERENCE]

89.2 Accept the model the market's own structure implies

Many authoritative sources, plus governed synchronization, plus provenance. [INFERENCE]

Many authoritative sources means writing down the authority map (Chapter 84, layer 4) instead of leaving it implicit, so that every system built on top knows which differences are breaks and which are legitimate divergence. A lender whose position differs from the agent's has a break. A lender whose accounting treatment differs from another lender's does not. **Systems that**

treat all differences as exceptions generate noise; systems that understand the authority map generate signal. *[INFERENCE]* This is the most immediately actionable insight in Part Ten and it requires no coordination at all to apply.

Governed synchronization means continuous rather than periodic reconciliation, with the recognition — from Chapter 74.4 — that a message stream carries events rather than state, and that completeness cannot be assumed. Detection is the cheap half; the utilities that stop at detection stop halfway (Chapter 44). *[INFERENCE]*

Provenance means each assertion travelling with the record specified in Chapter 71.1 and Chapter 88.3.

89.3 What this model gives up, honestly

It gives up elegance. It concedes that reconciliation is permanent rather than eliminable, that some duplication is structural, and that no participant will ever have a single unimpeachable view of a loan without asking someone else. *[INFERENCE]*

It gives that up because the alternative is not achievable and thirty years of evidence says so. **The distributed-ledger wave failed because it supplied consensus to a market that already had a contractually designated authority and needed propagation instead.** *[INFERENCE]* An industry that accepts the model above will build different things — and the things it builds have a chance of working.

What collaboration can and cannot solve

Part Seven established the consortium's capability and its limit. Stated as guidance for the parties who convene them:

90.1 Suitable for collaboration

Problems that are **shared infrastructure with symmetric benefit and no loser**: data pooling, settlement rails, messaging transport, connectivity, reference data. *[INFERENCE]* The evidence is that consortia perform well here — the bank-owned data platform reached material facility coverage; the trading venue took secondary-volume share from a standing start; the coexistence network reached hundreds of institutions with bank-reported improvements in fails and email traffic. *[FACT / VENDOR CLAIM]*

90.2 Unsuitable for collaboration

Anything that redistributes value between the members. *[INFERENCE]* Transparency is the paradigm case (Chapter 87.1). So is anything that would advantage the buy side at the sell side's expense, or the smaller agents at the larger agents' expense.

A consortium asked to solve a distributional problem will produce a product that solves the shared-infrastructure part and quietly omits the distributional part, and the omission will not be visible in the announcement. *[INFERENCE]* Anyone evaluating a consortium initiative should ask which half it is doing.

90.3 The honest note about consortium quality

The independent analysis quoted in Chapter 49 holds that a consortium product "doesn't need to be profitable; it just needs to preserve or strengthen the banks' existing competitive moats," and is "usually good enough, but not best-in-class." *[FACT]*

For the industry — as distinct from a competitor — that is a **feature with a cost**. The feature is that infrastructure gets built at all, funded by parties who can absorb its cost and who will not extract monopoly rents from each other. The cost is a quality ceiling, and a structural bias toward defending the layer rather than optimising it. *[INFERENCE]* An industry that relies exclusively on consortia gets adequate plumbing forever and never gets excellent plumbing. That is a defensible trade, and it should be made knowingly.

What one institution can do alone, and a realistic sequence

The paper's most practically useful finding is Chapter 41's third corollary: **an improvement that requires no other party to change escapes the mechanism that has defeated everything else.** *[INFERENCE]*

Everything below can be done by a single institution, this year, without an industry agreement, a consortium, a standard, or a regulator.

91.1 The unilateral list

Write down your own authority map. For each data element you hold, record who is authoritative, what your recourse is if you disagree, and which differences are breaks versus legitimate divergence. Even a basic authority map can improve reconciliation by distinguishing genuine breaks from differences that arise because two systems legitimately answer different questions. *[INFERENCE]*

Read your own agreements. Interpretation is newly economical (Chapter 62). An institution can extract the economically operative terms of its own portfolio — margin grids, prepayment triggers, PIK mechanics, waterfall priorities — into a structured representation with provenance to the clause, without asking anyone's permission, because they are its own documents.

[INFERENCE] This is the single largest capability change available and it requires nothing from the market.

Recompute independently. Having read the agreements, compute what the interest payment should be and compare it to what arrives. This produces earlier detection, evidence for disputes, and — the underrated one — **a control artefact for the institution's own audit and risk functions,** which is a defensible internal spend in a way that an efficiency project is not. *[INFERENCE]*

Measure your own break rate. No published loan-operations error or break rate exists anywhere in the public record (Appendix E). *[UNKNOWN]* An institution that measures its own has information nobody else in the market has, including about its own vendors.

Verify your own collateral. In asset-backed and receivables lending especially, establish independently that the assets exist and have been pledged once. The 2025–26 frauds reached named institutions through that gap. *[FACT]*

Fix the crossing, not the model. Require that every AI-derived term entering a system of record carries the provenance record of Chapter 88.3. This is a policy decision, not a technology purchase. *[INFERENCE]*

Buy agency service as a term. At syndication, treat the agent's operational standard as something to be negotiated rather than accepted (Chapter 86.3). *[INFERENCE]*

91.2 A realistic sequence

Not a technical architecture — an institutional progression, in the order the dependencies actually run:

1. **Interpretation.** Read your own documents. Requires: nothing external.
2. **Provenance.** Attach evidence to every extracted assertion. Requires: a policy and a schema.
3. **Independent verification.** Recompute positions, accruals and payments; compare. Requires: 1 and 2.
4. **Governed synchronization.** Reconcile continuously against counterparties, distinguishing breaks from legitimate divergence using the authority map. Requires: 3, and a counterparty willing to exchange data — the first step that needs anyone else.
5. **Better controls.** Convert the above into audit evidence and internal control. Requires: 3, and an audit function that will accept it.

6. Selective common infrastructure. Where several institutions have done 1 through 5 and discovered they are solving the same problem, pool it — through a consortium if the benefit is symmetric (Chapter 90.1), through a vendor if it is not.

7. Contractual reform. Change the agreement provisions that Chapter 39.1 identifies as the settlement bottleneck. Requires: origination-side agreement, and propagates at the speed of refinancing.

Three things about this sequence.

It is ordered by **coordination requirement**, which is the paper's binding constraint. Steps 1 through 3 require nobody. Step 4 requires one counterparty. Step 6 requires several. Step 7 requires the borrower.

It **does not end in replacing the servicing core**, and nothing in this paper's evidence suggests it should. An institution that completes all seven steps has a documented, verified, well-governed view of its own book and a better relationship with its counterparties' data. It still has the incumbent servicing platform, and Chapter 30 explains why that is the rational outcome rather than a failure of ambition. *[INFERENCE]*

And it is **available now**. Steps 1 through 3 use capability that did not exist in 2022 and does exist in 2026, applied to documents the institution already owns. *[INFERENCE]* That is the concrete, unglamorous form of the answer to this paper's central question.

The sequence implies a different modernization direction from the one the market has usually pursued.

Historically, many initiatives began with the market: establish a standard, create shared infrastructure, persuade participants to adopt it, and then improve the institution.

The sequence above begins inside the institution. Interpret what it already possesses. Preserve provenance. Verify independently. Improve controls. Synchronize selectively where another party is required. Pool

infrastructure only where repeated bilateral experience demonstrates a genuinely shared problem. *[INFERENCE]*

The first AI-native transformation of institutional credit may therefore happen inside institutions before it happens between them. *[INFERENCE]*

This is not a prediction that market infrastructure becomes unnecessary. It is a claim about sequence. For the first time, several meaningful improvements can begin before the coordination problem is solved.

FIGURE 16

What one institution controls, and what it does not



Interpretation, provenance, verification and controls sit inside the institution; reconciling against a counterparty needs one other party; shared standards and common infrastructure need the market. Chapter 91 orders the same material by coordination requirement, which is the paper's binding constraint. [INFERENCE]

The agenda in one place, and where each item comes from

Ten chapters is a long agenda, and an agenda that cannot be traced back to the analysis is a consulting deck. Every item below is generated by a specific finding in Parts One through Ten; the third column is the constraint that decides whether it happens.

THE PROPOSAL	THE FINDING THAT GENERATES IT	COORDINATION REQUIRED
Leave credit economics, covenant architecture and risk allocation bespoke. Standardize the container, negotiate the contents	Commercially significant terms are a negotiated price, not a data-quality defect (Ch. 36)	None — this is a prohibition, not a project
Standardize at the semantic layer: what a term <i>means</i> , versioned and governed	Six standardization layers have very different prospects; the semantic layer is the highest-value one never seriously attempted, and the Chapter 41 asymmetry is materially weaker there — not absent (Ch. 84)	Industry, through a trade body
Translate identifiers rather than replacing them	Identifier fragmentation is a competitive equilibrium among vendors who each own a scheme, not a coordination failure among users; the trade body's own answer was a cross-reference (Ch. 37.3, 85)	Bilateral to industry; already partly built
Standardize the machinery of transfer restrictions, not their substance	Consent regimes cause the settlement times, and the trade body's own diagnosis is that changing agreement provisions is necessary — but who holds the debt is a legitimate borrower interest (Ch. 22.3, 39.1, 83)	Origination-side; propagates at the speed of refinancing
Pay the party that bears the coordination cost — by fee, by ownership, by commercial incentive or by contractual reform	Initiatives fail when cost falls on the agent and benefit on the lenders, absent compensation or compulsion. Ownership is the one mechanism with a proven instance; fees are unexplored (Ch. 41, 86)	Varies by mechanism; ownership is proven
Accept that the agent stays central, and design around that rather than against it	The register is a contractual allocation of authority, not an accident of history. Every initiative that required the agent to surrender it has failed; the one that compensated the agent with ownership succeeded (Ch. 9, 41, 46, 69)	None — this is a design constraint
Adopt many authoritative sources with governed synchronization and provenance; stop pursuing a golden source	Authority is distributed by design and the design is contractual; what the market lacks is propagation with evidence, not agreement (Ch. 68–71, 89)	Bilateral, growing to industry
Bound AI to proposal, and keep proof deterministic	Interpretation is the one category that changed; calculation, coordination, synchronization and institutional control did not. No automated metric predicts extracted-specification quality (Ch. 26, 63, 65, 88)	None — an institutional policy
Reserve regulation for post-trade transparency	The parties who control the infrastructure lose from transparency, which is why thirty years of voluntary action has not produced it. This is the one problem the evidence says voluntary action cannot reach (Ch. 53, 87.1)	Regulatory

THE PROPOSAL	THE FINDING THAT GENERATES IT	COORDINATION REQUIRED
Do the unilateral work first: read your own documents, verify your own positions, know your own break rate	Steps 1 through 3 of the sequence require no other party's cooperation, use capability that did not exist in 2022, and apply to documents the institution already owns (Ch. 91)	None

Read down the last column and the paper's own prediction is visible. The items requiring no coordination are the ones that will happen. The items requiring an industry are the ones with a thirty-year record of not happening, and they do not become more likely because they are correct.

That is the uncomfortable shape of an industry agenda in a market with this structure, and it is the reason Part Thirteen has to be asked separately. A market can have a rational direction that no participant is individually incentivised to take, and no single firm can profitably push it there.

XIII

PART THIRTEEN

Paths for a new entrant

Part Twelve set out what the industry could change. This part asks a different question: **whether any single firm can profitably cause it to change.**

The distinction is not academic. Part Twelve's recommendations are collectively rational and, in several cases, individually unfundable — an improvement can be worth more than it costs while no party who could make it captures enough to justify making it. That is what Chapter 41's mechanism does to a market, and it is why an industry direction and a business opportunity are separate findings.

Six paths are set out below. Each is described in the same structure, so that they can be compared rather than advocated: the proposition; who buys it; what must be true; the structural advantage; the access dependency; the coordination dependency; the incumbent response; what kills it; the most likely failure mode; and the assessment.

Two of those deserve a word, because they are the ones that decide outcomes in this market and are usually left implicit.

Access dependency asks what the path needs from a party that does not have to give it. Chapter 74 established that this is a legal question before it is a technical one, and Chapter 81's verdict turned on it.

Coordination dependency asks how many parties must change behaviour for the path to work. Chapter 41 established that this is the market's characteristic point of failure, and Chapter 48 found that across six

technology waves the base rate for initiatives requiring more than one party to change is zero.

A path that needs neither is not thereby a good business. But a path that needs both is, on this paper's evidence, not a business at all.

None of these paths is "build a better Loan IQ." Part Four established that this competes on four dimensions where product quality is largely beside the point, and nothing in Parts Five through Twelve weakened that finding.

Path A: the governed coexistence layer

The proposition

Sit alongside the incumbent servicing core at an agent bank. Receive the message and notice flow. Read the underlying agreements with AI. Maintain a parallel model of positions, accruals and cash. Reconcile continuously against the incumbent. Surface breaks with provenance to the clause. Optionally, over time, take over functions.

Who buys it

Ambiguous, which is the first warning sign. The efficiency case addresses operations, which has the smallest discretionary budget in the institution (Chapter 8). The control case addresses risk and audit, which is a better buyer but wants a different product. *[INFERENCE]*

What must be true

H1 through H5 of Chapter 75.4. In practice: the institution grants observation access; the reconstruction is complete enough; the accuracy is verifiable; and someone with a budget wants it.

Structural advantage

It asks the institution to replace nothing, which means it faces none of Layers 2 through 5 of the moat (Chapter 30.1). It is architecturallyprecedented — event interception, anti-corruption layer, parallel run are named patterns with published risk profiles (Chapter 75.2). And it addresses a cost the institution genuinely bears. *[INFERENCE]*

Access dependency

Total, and unprecedented. The path requires an institution to permit a third party to receive its loan servicing traffic. No documented instance exists, the decision belongs to security, vendor risk and the general counsel

rather than to the operations team that wants the capability, and the regulatory trajectory in two jurisdictions runs from interception toward sanctioned interfaces (Chapter 74). *[FACT / UNKNOWN]*

Coordination dependency

Low. Once access exists, the path is sold to one institution for that institution's benefit and escapes the Chapter 41 mechanism entirely. The dependency is concentrated in a single gate rather than spread across a market. *[INFERENCE]*

Incumbent response

The incumbent already sells the translation layer, ships interface products, and can price or bundle against a translation-based differentiator (Chapter 74.5). *[FACT / VENDOR CLAIM]* The bank-owned network in the same position has the cap table and the servicing-core integration (Chapter 77.1). Success here is also the trigger condition for a consortium response (Chapter 50). *[INFERENCE]*

What kills it

The access gate (Chapter 74). The occupied position (Chapter 77.1). The self-defeat problem: the version that clears change management prices as a reporting tool (Chapter 77.5). And a translation layer the incumbent already sells (Chapter 74.5).

Most likely failure mode

Not refusal — **dilution**. The version that survives security review, vendor risk and change management is the read-only, report-generating version, and that version prices as a reporting tool rather than as infrastructure. The path does not die at the gate; it gets through the gate having become something smaller. *[INFERENCE]*

Assessment

The most intellectually satisfying path and the one with the weakest evidentiary support. *[INFERENCE]* It is architecturally sound, preceded in banking generally, and squarely in the path of AccessFintech — which has the bank cap table and the Finastra integration.

Two qualifications keep this from being a dismissal. First, the finding against it rests substantially on an **absence of precedent** for observation access, not on a demonstration that access cannot be obtained (Chapter 81.4). An entrant who secures it has not defied a structural law; it has done something for which this paper found no prior instance. Second, AccessFintech's identified limitation — that it reconciles data rather than reading the underlying agreements — is real, even if reading agreements is now table stakes elsewhere. *[INFERENCE]*

Recommendation: not a first move. A reasonable second move for a firm that has already established access and credibility through another path. *[INFERENCE]*

Path B: the private-credit administration platform

The proposition

Perform the administration function for private-credit managers — register, servicing, notices, cash, covenant monitoring, investor reporting — as a technology-enabled service rather than as software sold to someone else to operate.

Who buys it

Private-credit managers between roughly \$1 billion and \$25 billion of assets: large enough that spreadsheets have broken, too small to have built an operations function, and with an existing budget line because they are already paying an administrator. *[INFERENCE]*

What must be true

That the service can be delivered at a cost structure meaningfully better than incumbent administrators, and that managers will switch administrators — which, unlike banks switching servicing cores, they do routinely. *[INFERENCE]*

Structural advantage

The access problem dissolves rather than being solved. A party performing the administration function holds the data by right, because it is operating its own system rather than observing someone else's. It also inherits an existing budget line, since the manager is already paying an administrator, and it operates in the one segment where the operating-layer evidence is corroborated by a supervisor rather than only by vendors (Chapters 57, 60). *[INFERENCE]*

Access dependency

None. This is the path's defining property and the reason it appears in Chapter 80 as the only clean solution to the obstacle the coexistence hypothesis treats as central. *[INFERENCE]*

Coordination dependency

None. One manager decides; nobody else has to change anything. *[INFERENCE]*

Incumbent response

The entrenched administrators can adopt the same tooling and compete on price, and several of them own their own technology and would be defending a platform as well as a book of business (Chapter 11). Their advantage is scale, brand and the fact that switching administrators — while routine — still carries operational risk for the manager. *[INFERENCE]*

What kills it

Entrenched administrators with scale, brand and existing relationships, several of which own their own software (Chapter 11). Services economics rather than software economics. And the possibility that the incumbents adopt the same tooling and compete on price. *[INFERENCE]*

Most likely failure mode

Margin compression before differentiation. The path scales with people until the tooling is genuinely better, and the window in which it is genuinely better is the window before incumbent administrators buy or build the equivalent. A firm that grows headcount faster than it grows differentiation ends up as a smaller version of the competitor it set out to displace. *[INFERENCE]*

Assessment

The path that most cleanly dissolves the paper's central obstacle.

[INFERENCE] The access problem of Chapter 74 disappears entirely, because a party performing the administration function holds the data by right rather than by permission — it is not observing someone else's system, it is operating its own. The buyer exists and has a budget. The evidence in Chapter 57 that private credit's operating layer is spreadsheets is the strongest problem-evidence in the paper, and it is corroborated by a supervisor rather than only by vendors.

Its weakness is the honest one: it is a services business, it scales with people until the tooling is genuinely differentiated, and the multiple is lower. That is a real cost. It is not a reason to dismiss the path — several of the market's most successful firms in this layer are services businesses that own their technology (Chapter 11) — but it is a different company from the one most entrants intend to build. *[INFERENCE]*

Path C: the verification utility

Chapter 69 identified three rows of the authority table where no party is authoritative. Chapter 81.3 established that they are not one opportunity. This path is two propositions, and they should be assessed separately.

94.1 Collateral integrity

The analysis sits in Chapter 73 and Appendix F and is not repeated here. What this path needs from it is four findings: the filing system establishes rank rather than existence and was designed to; an authoritative verification mechanism existed and was unused in every case examined, so most of the gap is operational; the one part no single institution can close alone is whether an asset sits in another lender's facility; and the mechanism that would close *that* requires other institutions to participate, which places it under Chapter 41.

The proposition. Establish independently that a claimed asset exists, that it has been pledged once rather than twice, that a borrowing base is what it is represented to be, and that a covenant calculation follows from the agreement and the financials as presented — continuously rather than periodically, and from evidence the borrower does not control.

Who buys it. Lenders after a loss, and lenders whose credit committees have decided not to have one. Potentially also the trustees and collateral agents who bear responsibility for the underlying assets, and the insurers behind them. *[INFERENCE]*

What must be true. That there is a verifiable external fact — and here there is. *[INFERENCE]* That verification can be performed from borrower-side data, registry data, counterparty confirmation and document analysis without requiring access to a servicing core (test T7, Chapter 79). *[HYPOTHESIS]* And

that institutions will pay for a capability most of which they could have exercised themselves and did not — which is a question about procurement behaviour rather than about the problem. *[UNKNOWN]*

What kills it. If verification requires the same access as Path A, it inherits Path A's gate. If the answer to the preceding question is that the market treats this as a diligence failure to be fixed by rewriting audit rights rather than by buying a system, the buyer does not appear — and the post-failure reforms reported so far are of exactly that kind. *[INFERENCE]* If market memory fades before a product exists, the urgency closes. And Setpoint already works the adjacent problem in asset-backed lending, performing collateral document verification, borrowing-base computation and compliance monitoring at more than \$50 billion of financed assets a year.

[VENDOR CLAIM]

Structural advantage. It contests no existing authority, requires no observation access to a servicing core, needs no takeover path, has a verifiable external fact at its centre, and has a loss-prevention buyer — the least discretionary spending category in a credit institution after regulatory compliance. It is also the one proposition in the paper that survives the failure of H3 and H5 (Chapter 81.5, test T7). *[INFERENCE]*

Access dependency. Borrower-side, not incumbent-side — collateral schedules, borrowing-base data, invoices and registry queries, obtained under a lending relationship that already entitles the lender to them. That is a materially easier ask than Path A's, and it is a hypothesis rather than a demonstrated fact: test T7 of Chapter 79 exists precisely to settle whether the work can be done without servicing-core access. *[HYPOTHESIS]*

Coordination dependency. Split, and the split is the whole assessment. Existence, eligibility and borrower-reporting verification require no other institution's agreement and can be sold one lender at a time. Cross-facility uniqueness — the part that is genuinely structural — requires the other

lenders, and therefore inherits Chapter 41 and the non-linear adoption problem of §73.8: partial coverage yields a warning signal, only complete coverage yields a proof. *[INFERENCE]*

Incumbent response. An existing venture already works the adjacent problem in asset-backed lending at scale *[VENDOR CLAIM]*, and the post-failure reforms reported so far are contractual — expanded audit and confirmation rights, direct obligor confirmation, unscheduled receivable-level audits — which is the market solving the problem by rewriting agreements rather than by buying a system. That is a competitive response of a kind a software entrant cannot outbid. *[INFERENCE]*

Most likely failure mode. The problem is conceded and the budget never appears. Every prospect agrees the failure mode is real, agrees the verification is worth doing, and concludes that it is their own diligence to fix rather than a system to procure — which is exactly the T4 failure mode, arriving in the one place where the problem evidence is strongest. *[INFERENCE]*

Assessment. The best-evidenced proposition in the paper, which is a narrower claim than "the strongest opportunity." *[INFERENCE]* Its strength is structural: it contests no one's authority, requires no one else to change for the operational part, has a verifiable fact at its centre, has a loss-prevention buyer, and has a recent failure at named institutions as its proof. Its weaknesses are two and they are specific. The part with the clearest buyer is the part the institutions could already have done, which makes it a services-and-discipline sale as much as a software one. And the part that is genuinely structural — uniqueness across facilities — is the part that needs an industry, and therefore inherits the base rate of everything else in this paper. A demonstrated problem is not a demonstrated business.

94.2 Private-credit valuation governance

The proposition. Provide the process, evidence and auditability around manager-determined fair value that the supervisors have named as missing.

Who buys it. This is the problem. The manager produces the valuation; the manager would be the buyer; and the buyer's interest in being examined is structurally weaker than its interest in being able to say it has been examined. The parties with the strongest interest in the answer — limited partners and supervisors — do not buy software. *[INFERENCE]*

What must be true. That compliance pressure converts into procurement. The supervisory tailwind is real and named across three institutions in two jurisdictions. *[FACT]* Whether it produces budget is a separate question the paper cannot answer. *[UNKNOWN]*

What kills it. The conflicted buyer. The absence of an external fact to verify against — "verification" here means auditing a process, which competes with the audit and valuation-advisory firms that already do it. And the possibility that supervisors settle for disclosure rather than process evidence. *[INFERENCE]*

Assessment. A genuine problem with an awkward buyer. *[INFERENCE]* It should not be bundled with collateral integrity, as an earlier draft of this paper did: the two share only the property that no external authority exists, and differ on the question that determines whether a business is available — whether there is a fact of the matter.

Path D: the channel play

The proposition

Build the capability and take it to market through the implementation boutiques, systems integrators and managed-service firms already inside these institutions.

Who buys it

The channel partner buys it as something to sell; the institution buys it as part of a programme it already trusts. *[INFERENCE]*

What must be true

That the capability expands the partner's billable footprint rather than substituting for it (Chapter 13). The boutiques already resell partner technology into their installed base, which is a working template. *[FACT]*

Structural advantage

It borrows a relationship rather than building one, which is the scarcest asset in this market. Chapter 13 established that the ecosystem's economics reward anything additive and punish anything replacing — so a coexistence product is, for the boutiques and pure integrators, something new to implement, test and support on an installed base they already own. *[INFERENCE]*

Access dependency

Indirect but real. The partner has the institutional relationship; it does not thereby have the right to grant a third party access to servicing traffic. A channel shortens the sales cycle and does not open the gate of Chapter 74. For an additive product that needs no such access, the dependency is low; for Path A delivered through a channel, it is unchanged. *[INFERENCE]*

Coordination dependency

Low, but with an added party. One institution still decides, and now a partner has to want the sale. That is a commercial dependency rather than a market-coordination one, and it is a different failure surface from Chapter 41's. *[INFERENCE]*

Incumbent response

The managed-service firms that own their own servicing technology are competitors in either direction (Chapter 13.6), and the global integrators sell competing modernization narratives of their own. The channel is not neutral territory. *[INFERENCE]*

What kills it

Margin capture and relationship capture by the channel — the partner owns the account, sets the price, and can substitute a competitor's product without the institution noticing. Slow, indirect feedback from end users, which is a product-development problem as much as a commercial one. And the global integrators are frenemies: distribution at scale, but competing modernization narratives of their own. *[INFERENCE]*

Most likely failure mode

Becoming a line item in someone else's programme. The product ships, the partner bills the hours, the institution associates the outcome with the partner, and the vendor is left with component margins and no direct relationship to the user whose feedback it needs. *[INFERENCE]*

Assessment

Not a standalone path but the best available answer to the anchor problem. *[INFERENCE]* Chapter 51 established that the first anchor is brutally slow, that nobody wants to be first, and that the anchor knows its signature is the asset and will price accordingly. A channel partner with an existing trusted relationship compresses that, at the cost of margin and of direct

customer contact. For a firm without institutional relationships, that trade is usually worth making; for a firm that has them, it is usually not. *[INFERENCE]*

Path E: sell to the incumbents

The proposition

Build the capability and sell or license it to the parties who need it and cannot easily build it: Finastra, with a large legacy core and dispersing implementation expertise (Chapter 34); administrators such as Alter Domus and SS&C defending services margins; and data incumbents such as S&P Global and Morningstar assembling private-credit capability.

Who buys it

The evidence is that these buyers are active: five acquisitions of extraction-adjacent companies since 2021, including Morningstar's purchase of Lumonic and Alter Domus's of Solvas. *[FACT]*

What must be true

That the capability is genuinely hard to replicate and that the buyer's build-versus-buy calculus favours buying.

Structural advantage

The buyers are demonstrably active, the sale requires no institutional access gate, no anchor customer and no market coordination, and the acquirer supplies the distribution the entrant lacks. It is the only path in this list whose success condition has recent, repeated, documented instances. *[FACT / INFERENCE]*

Access dependency

None from an institution. The dependency is on a corporate counterparty's build-versus-buy judgement, which is a negotiation rather than a permission. *[INFERENCE]*

Coordination dependency

None. One buyer decides. *[INFERENCE]*

Incumbent response

This path is the incumbent response, purchased. The risk is not that an incumbent competes but that it waits, builds, or buys a competitor — and that the entrant's leverage decays with every month the capability becomes less rare. Extraction is already table stakes and consolidating (Chapter 64). *[INFERENCE]*

What kills it

Single-buyer dependency in a consolidating market — the lesson of Chapter 52, where a venture's fate was determined by a transaction it was not party to. And supplier economics: a component sold into someone else's platform earns component margins.

Most likely failure mode

The window closes quietly. The capability stops being scarce, the acquisition multiple compresses, and a firm that built for a sale finds itself running a component business it did not intend to run. *[INFERENCE]*

Assessment

The highest-probability, lowest-ceiling path, and the one the market's transaction record most supports. *[INFERENCE]* It is the correct answer for a firm optimising for a defined outcome rather than a market position, and the wrong answer for a firm that wants to own a layer.

Path F: decline

The proposition

Conclude that the market's structural defences exceed the available advantage, and do something else.

The case

It is set out in full in Chapter 77 and it is not weak. The coexistence position is occupied by a bank-backed network with a decade of network effects. The access gate has no documented precedent. The consortium wall is extending into the target segment with a mega-manager inside it. Extraction is table stakes and consolidating. Three decades of evidence say the market adopts what requires no change, and what requires no change may not be valuable enough. The incumbent is modernising and re-aiming rather than vacating. And the sales cycle is six to twelve months at best for an anchor that co-owns the competition. *[FACT / INFERENCE]*

What would have to be true for this to be wrong

Any one of four things, and a reader who believes one of them should not decline. That observation access is obtainable and this paper simply could not find the instance (Chapter 81.4). That private credit resists consortium formation because there is no agent chokepoint to pool (Chapter 60.2). That the collateral proposition converts into procurement rather than into contractual reform (test T4). Or that a downturn's urgency effect dominates its budget-freeze effect (Chapter 59.4). *[HYPOTHESIS]*

Most likely failure mode

This path's failure mode is not commercial. It is **declining a market that was, on evidence not yet public, open** — which is unobservable to the party that declined, and is the reason the paper states the case for declining at full strength rather than dismissing it. *[INFERENCE]*

Assessment

A defensible conclusion that this paper does not reach but does not dismiss. *[INFERENCE]* A reader who weights the access gate and the consortium clock more heavily than this paper does should reach it. The disagreement is locatable and legitimate, and Appendix C says exactly where it sits.

The six paths, compared

Before the tests, the comparison. Each column is a constraint the paper established rather than a criterion invented here, and the point of setting them side by side is that no path is strong on all of them.

	A COEXISTENCE LAYER	B PC ADMINISTRATION	C VERIFICATION UTILITY	D CHANNEL	E SELL TO INCUMBENTS	F DECLINE
Buyer exists with a budget	Ambiguous	Yes — already paying an administrator	Yes after a loss; unproven before one	The partner, then the institution	Yes — demonstrably active	n/a
Access dependency	Total, unprecedented	None — holds data by right	Borrower-side; hypothesis T7	Indirect; unchanged for A	None from an institution	n/a
Coordination dependency	Low, once access exists	None	Split: none for existence, all lenders for uniqueness	Low, plus a partner	None	n/a
Contests an existing authority	No	No	No	No	No	n/a
Incumbent response	Sells the translation layer; occupied position	Administrators adopt the tooling	Contractual reform instead of procurement	Frenemy integrators; platform- owning BPOs	Waits, builds, or buys a rival	n/a
Economics	Platform, if it survives dilution	Services, improving with tooling	Feature-sized until it is a platform	Component, channel takes margin	Component; defined outcome	n/a
Most likely failure mode	Dilution at change management	Margin compression before differentiation	Problem conceded, budget never appears	A line item in someone else's programme	The window closes quietly	Declining an open market
Rests on an absence of evidence	Yes — H3, H5	No	Partly — T7	No	No	Yes, in the other direction

[INFERENCE]

Two patterns are visible reading across rather than down.

None of the six is defeated by contesting an authority, because none of them tries to. That is the one constraint the paper's own framework says is fatal, and every path here has already been shaped around it. The differences are all in access, coordination and buyer.

The paths that dissolve the access problem do so by giving up platform economics, and the paths that keep platform economics depend on an access grant nobody has been documented to give. That trade is the honest summary of the entrant's position in this market, and it is not resolved by any of the six. *[INFERENCE]*

Decision criteria independent of path

1 Which complexity category does the value come from? <i>CH. 26</i>	2 Does anyone else have to change behaviour? <i>CH. 48</i>
3 Who bears the cost and who gets the benefit? <i>CH. 41</i>	4 Does this contest an existing authority? <i>CH. 69</i>
5 What is the access mechanism, and has anyone been granted it? <i>CH. 74</i>	6 Who is the buyer, and does it have a budget line today? <i>CH. 79</i>
7 Is the differentiator already sold or bought by an incumbent? <i>CH. 64, 74.5</i>	8 What happens when it works? <i>CH. 50</i>
9 What is the single-partner dependency? <i>CH. 52</i>	10 How is the model's output verified? <i>CH. 63</i>

Whatever path is chosen, the analysis in this paper implies a set of tests that should be applied before commitment. They are not a checklist of good practice. Each one is the compressed form of a finding established earlier, each has a specific failing answer, and each has killed something real in this market's history.

They are stated as questions, with the chapter that generates the test and the answer that should stop the plan.

1. Which complexity category does the value come from? (*Chapter 26*) The market's work decomposes into calculation, interpretation, coordination, synchronization and institutional control. Every technology that succeeded here addressed calculation or synchronization; every one that failed was aimed at coordination or institutional control; and what changed in 2026 was interpretation and nothing else. **Failing answer: Category 3 or 5.** A proposition whose value comes from getting parties to coordinate, or from reallocating who decides, is aimed at the two categories that have defeated every initiative surveyed here.

2. Does anyone else have to change behaviour for this to work? (Chapter 48)

Across six technology waves, the initiatives that were adopted asked nothing of anyone; the initiatives that required a second party to change were adopted partially or not at all. **Failing answer: more than one party.** The base rate of success in that condition, across the waves examined here, is zero.

3. Who bears the cost and who gets the benefit? (Chapter 41) The market's characteristic failure is not technical. It is that the party who must do the work — usually the agent — is not the party who receives the value. **Failing answer: different parties, cost-bearer uncompensated.** Chapter 86 sets out the four non-regulatory ways to compensate them, of which ownership is the only one with a proven instance and fees are the one nobody has tried.

4. Does this contest an existing authority? (Chapter 69) Authority in this market is allocated by contract, fact by fact, in advance. **Failing answer: the plan requires being right against a party that has been designated to decide.** Contesting the agent's register is contesting a contractual allocation and will not be won on the merits. Operating where no authority exists is the strongest available position — and Chapter 94 shows that even there, the presence or absence of a verifiable external fact decides whether a business exists at all.

5. What is the access mechanism, and has anyone ever been granted it? (Chapter 74) Access to another institution's data is a legal question before it is a technical one, decided by security, vendor risk and the general counsel rather than by the team that wants the capability. **Failing answer: no documented instance, and no route that does not run through interception.** That is then the highest-risk assumption in the plan, and it should be tested before anything is built — which is why it is test T1 rather than T7.

6. Who is the buyer, and does that buyer have a budget line today? (Chapter 79, T4) Operations has the smallest discretionary budget in the institution; control functions have smaller but far more defensible ones; loss

prevention buys after an incident and stops buying when the memory fades. **Failing answer: every prospect agrees the problem is real and none has a budget line.** This is the most common failure mode for products of this class, and Chapter 94.1 shows it arriving even where the problem evidence is strongest.

7. Is the differentiator something an incumbent already sells or has already bought? (Chapters 64, 74.5) Failing answer: yes, and it is the core of the pitch. Reading contracts is table stakes and consolidating; provenance is already being marketed under another name; translation to the incumbent core is sold by the incumbent, which can therefore price or bundle against it.

8. What happens when it works? (Chapter 50) The documented rule in this market is that banks compete on origination and cooperate on infrastructure, and that a layer becoming valuable is the trigger for pooling it. **Failing answer: no answer to "and then the banks build it together."** Success is not the end of the competitive problem; it is the start of it.

9. What is the single-partner dependency? (Chapter 52) Failing answer: one partner, one distribution route, or one acquirer, in a market consolidating at this rate. A venture's fate can be decided by a transaction it is not party to, and the paper's canonical instance won its litigation and filed for bankruptcy anyway.

10. How is the model's output verified? (Chapter 63) Failing answer: "the model is accurate." That is not an answer. No automated metric strongly predicts extracted-specification quality, human evaluation remains necessary, and a wrong margin grid produces no error — it produces a number. Any plan that depends on machine interpretation needs a verification design, and the cost of that design is a permanent line rather than a transitional one.

A note on how to use them. These are not weighted and they do not produce a score. Nine acceptable answers and one failing answer is a failing result, because the failures in this market have been single-point rather than cumulative: Symbiont failed on nine, the ledger wave failed on one, and the reason both are in this paper is that each failure was sufficient. *[INFERENCE]*

Closing

Three decades of capable people have tried to fix this market, and the pattern of their failure is consistent enough to be predictive.

They failed because they treated a problem of authority as a problem of data, a problem of incentive as a problem of technology, and a problem of contract as a problem of software. The distributed ledgers supplied consensus to a market that already had a designated authority. The messaging standards supplied syntax to a market whose problem was semantics. The settlement reforms supplied conventions to a market whose bottleneck was consent. Each was technically competent and aimed at the wrong layer.

What changed in 2026 is narrow and genuine. Machines can now read the documents this market is made of. That was not true four years ago, it was the market's binding constraint, and it is worth a great deal to whoever can convert it into something an institution will use.

What did not change is everything else. The agent still holds the register. The agreement is still negotiated to be unique. The consent still has to be given by someone with no reason to hurry. And the handful of institutions that own this market's infrastructure will still, when a layer becomes valuable, build it themselves and own it together.

For the industry, the implication is a smaller and more achievable agenda than the one usually proposed. Stop pursuing a golden source that the market's own contracts forbid. Standardize meaning rather than format, and provenance rather than identifiers — accepting that the semantic layer is cheaper to agree than the ones already tried, not that it is free. Translate where incumbents have already created competing standards, and allocate where the question is authority. Pay the party that bears the cost. Reserve regulation for the one problem — transparency — that the evidence says voluntary action cannot reach. And recognise that the most valuable work

available requires no industry agreement at all: an institution can read its own documents, verify its own positions, and know its own break rate, this year, alone.

For an entrant, the implication is narrower still. It is not in replacing the core, for which no demonstrated path exists and where four of the five defensive layers are indifferent to how good the product is. It is not, on the current evidence, in observing the core, which requires a permission nobody has been documented to grant. It is in the parts of this market where **no party is authoritative** — and, among those, specifically in the parts where there is nonetheless a fact of the matter: whether an asset exists, and whether it has been pledged once.

The market's own recent history is the argument, and it has to be stated precisely, because the loose version of it is wrong. In 2025 and 2026 two companies failed on collateral that had been counted twice or had never existed, and the exposure reached institutions that are very good at credit. The law was not defective: the filing system allocated rights exactly as designed, and it was never asked to establish that anything existed. Nor were the institutions helpless: in every case an authoritative source — a state title record, a customer who could confirm an invoice — was available and was not consulted. What no institution could do alone was see into another lender's facility, and that is the one place where nobody in the structure was authoritative and nobody was wrong.

The six paths are commercial consequences of the architecture, not the architecture itself.

The larger finding is simpler.

Institutional credit did modernize. It repeatedly improved documentation, data, workflow, reconciliation and infrastructure. The hardest problems persisted where improvement required another institution to change,

where authority was contractually allocated, where consent carried economic value, or where the party bearing the cost did not capture the benefit. *[INFERENCE]*

AI changes one important variable in that history. Bespoke financial language can now be interpreted at a scale and cost that was not previously practical. That creates new possibilities for institutional understanding, independent verification, provenance, knowledge preservation and control.

It does not grant authority.

It does not create consent.

It does not align incentives.

It does not make probabilistic output financially correct.

And it does not erase the systems and institutions that currently hold operative state.

The strategic opportunity therefore lies neither in assuming that AI changes everything nor in concluding that the surviving constraints mean it changes little.

AI changed what an institution can economically know. It did not automatically change what that institution is entitled to decide.

That distinction may determine where the next modernization wave begins.

It may begin inside institutions: with their own agreements, their own positions, their own evidence, their own controls and their own institutional knowledge.

Only then does the question become how better-governed institutional states synchronize with one another.

Technology can increasingly interpret what was previously trapped in bespoke documents and scarce human expertise.


It still cannot grant authority.

Any architecture that ignores that distinction risks repeating thirty years of modernization history with better technology.

That is the shape of the opportunity, and it is also the shape of the problem. This market does not fail because its participants are careless. It fails at the joints between them, where nobody's authority reaches and everybody's assumption does.

How this research was produced

This paper was produced through a governed institutional research process developed at Conektiq.



The system was used to research the market, preserve institutional context, distinguish fact from inference, identify contradictions, challenge proposed conclusions, and maintain the evidence behind the final analysis.

The technology did not supply the point of view. It made it possible to investigate, challenge, and refine that point of view at a scale that would previously have required a substantial research team.

The paper's evidence classifications — FACT, VENDOR CLAIM, INFERENCE, HYPOTHESIS and UNKNOWN — are part of that process.

APPENDIX A

Chronology, 1970–2026

Dates are given as reported. Where a source's precision could not be confirmed, the entry is marked (*approx.*).

The formation period

YEAR	EVENT
1970	Misys founded as an insurance-software provider; later expands into banking software via acquisition. <i>[FACT]</i>
1970s	Midas core-banking system developed — a separate product line frequently and wrongly cited as Loan IQ's ancestor. <i>[FACT]</i>
1981	Delaware's Financial Center Development Act eliminates interest-rate ceilings on revolving and closed-end credit and offers declining bank tax rates, requiring out-of-state holding companies to employ at least 100 people within a year. <i>[FACT]</i>
1980s	Leveraged buyout boom creates a class of highly leveraged corporate loans attractive to non-bank investors.
Late 1980s–early 1990s	Distressed bank debt trading emerges out of necessity; a rudimentary secondary market forms.
1991	Delaware hosts 31 credit-card and wholesale banks with roughly 13,000 direct and 7,000 related jobs. <i>[FACT]</i>
1991	ACBS design team forms. <i>[VENDOR CLAIM]</i>
1993	ACBS first customer. <i>[VENDOR CLAIM]</i>
1995	LSTA founded , New York. <i>[FACT]</i>
1996 (Dec)	LMA founded , London. <i>[FACT]</i>
1998	APLMA founded , Hong Kong, by fifteen international banks. <i>[FACT]</i>
1998	Delayed compensation introduced as the market's remedy for late settlement. <i>[FACT]</i>
1999	Loan mark-to-market pricing service positioned as the leading source of objective loan prices. <i>[VENDOR CLAIM]</i>
2000	A major investment bank establishes its US back office in Salt Lake City. <i>[FACT]</i>

Consolidation and the first modernization waves

YEAR	EVENT
2003 (Jan–Apr)	Fidelity National Financial acquires ALLTEL Information Services' financial-services division for ~\$1.05bn, announced 29 January, completed 1 April; renamed Fidelity Information Services. ACBS enters the FIS family here — not via SunGard. <i>[FACT]</i>
2004 (Jan)	Misys acquires IQ Financial Systems from Deutsche Bank for \$40m cash. IQFS had ~\$22m FY2003 revenue and ~29 clients. Loan IQ's origin. <i>[FACT]</i>
2005	SunGard taken private for \$11.4bn by a seven-firm consortium at \$36.00 per share — one of the largest technology buyouts ever. <i>[FACT]</i>
2005	TPG and Thomas H. Lee Partners take a large stake in FIS. <i>[FACT — APPROXIMATE PERCENTAGES]</i>
2006 (Dec)	SunGard acquires Automated Financial Systems <i>Corporation</i> , a brokerage/wealth vendor — not the Exton, PA commercial-lending firm of a similar name. <i>[FACT]</i>
2008 (Jul–Sep)	Markit acquires FCS Corporation — Wall Street Office — from JPMorgan. Combined footprint cited at 400+ institutions and over \$1tn of assets; price undisclosed. <i>[FACT]</i>
2008	DTCC white paper calls for automation of syndicated loan processing; notes the industry "continues to rely on faxes as its primary communication method." <i>[FACT]</i>
2008	DTCC Loan/SERV position-reconciliation service launched. <i>[FACT]</i>
2009 (Oct)	Markit acquires ClearPar from Fidelity Information Services. <i>[FACT]</i>
2010	Markit acquires the DTCC Loan/SERV messaging portal and folds it into WSOData. Contemporaneous reporting: WSOData receives >7 million agent notices a year and redistributes >4 million. <i>[FACT]</i>
2011	A settlement vendor builds an FpML-based matching product; it "has yet to win any users." <i>[FACT]</i>
2012 (Mar)	Vista Equity Partners takes Misys private for ~£1.3bn. <i>[FACT]</i>
2013	Interagency Leveraged Lending Guidance issued by the Federal Reserve, OCC and FDIC. <i>[FACT]</i>
2015 (Aug)	FIS agrees to acquire SunGard for ~\$9.1bn including assumed debt. <i>[FACT]</i>
2016 (Mar–Jul)	Markit and IHS merge in an all-stock deal valued at more than \$13bn, forming IHS Markit. <i>[FACT]</i>
2016 (Sep 1)	LSTA delayed compensation moves from no-fault to requirements-based. Q2 2016: only 23% of par trades settled within T+7. <i>[FACT]</i>
2016	AccessFintech founded. <i>[FACT]</i>

The DLT wave and the consortium era

YEAR	EVENT
2017 (Mar–Jun)	Vista acquires D+H at C\$25.50/share, ~C\$4.8bn enterprise value; Finastra formed June 2017 from Misys + D+H. <i>[FACT]</i>
2017	GAO rules the 2013 leveraged-lending guidance was a "rule" under the Congressional Review Act never submitted to Congress. <i>[FACT]</i>
2018 (Feb 9)	DC Circuit holds open-market CLO managers are not "securitizers" — exempt from 5% risk retention. Unlocks US CLO growth. <i>[FACT]</i>
2018 (Apr)	Finastra's Fusion LenderComm goes live on an enterprise ledger platform — first application live on that platform; ~7 pilot banks. <i>[FACT]</i>
2018 (Mar)	Deutsche Bank inadvertently transfers €28bn (~\$35bn) to an exchange account; caught within minutes. <i>[FACT]</i>
2018 (Oct 17)	A global consultancy acquires a New York Loan IQ implementation specialist. <i>[FACT]</i>
2019	Further banks sign to LenderComm. No public reporting on its blockchain deployment after ~2019. <i>[FACT/UNKNOWN]</i>
2019 (Sep)	Vista forms Allvue by acquiring AltaReturn and merging it with Black Mountain Systems. <i>[FACT]</i>
2019 (Nov)	The syndicated-loan DLT joint venture's CEO resigns; a Delaware court orders it wound down. <i>[FACT]</i>
2019	JPMorgan reported weighing a shift of thousands of New York-area back- and middle-office roles to lower-cost hubs. <i>[FACT]</i>
2020 (Aug 11)	Citibank wires ~\$893.9m of Revlon principal instead of ~\$7.8m of interest. Six Delaware staff and nine offshore contractors in the control chain. <i>[FACT]</i>
2020 (Oct 7)	OCC assesses a \$400m civil money penalty against Citibank for risk-management, data-governance and internal-control deficiencies. <i>[FACT]</i>
2020 (Nov)	Only 17.7% of one trade body's members using or considering blockchain, against 60.4% for plain electronic platforms. <i>[FACT]</i>
2020 (Oct)	AccessFintech Series B, \$20m. <i>[FACT]</i>
2021 (Feb 16)	SDNY rules Revlon lenders may keep ~\$500m under the discharge-for-value doctrine. <i>[FACT]</i>
2021 (Aug)	Finastra publishes "Busting the Loan IQ myth." <i>[FACT]</i>
2021 (Oct)	Finastra's LenderComm integrated with AccessFintech's data layer, with updates flowing back to Loan IQ. <i>[FACT]</i>
2021 (Oct 22)	European Commission conditionally clears S&P Global / IHS Markit, requiring divestiture of the leveraged-loan data business. <i>[FACT]</i>

YEAR	EVENT
2021 (Nov 1)	A registrar group completes the ~\$750m acquisition of a major bank's US corporate-trust services business. <i>[FACT]</i>
2021 (Nov 12)	DOJ requires substantial divestitures and a non-compete waiver to clear S&P Global / IHS Markit — at the price-reporting layer, not at settlement or administration. <i>[FACT]</i>
2021 (Sep)	Allvue shelves its IPO after seeking a ~\$1.7bn valuation. <i>[FACT]</i>
2021 (end)	\$53m settlement paid to the DLT venture. <i>[FACT]</i>

The current period

YEAR	EVENT
2022 (Feb 28)	S&P Global completes its merger with IHS Markit , ~\$44bn enterprise value. Loan pricing, identifiers, settlement and administration under one owner. <i>[FACT]</i>
2022 (Mar 16)	Versana announced by JPMorgan, Bank of America, Citi and Credit Suisse. <i>[FACT]</i>
2022 (Apr)	S&P Global agrees to sell Leveraged Commentary & Data and the leveraged-loan indices to a competitor for up to \$650m. <i>[FACT]</i>
2022 (Jun 14)	Octaura announced , backed by seven dealer banks and a ratings and analytics firm. <i>[FACT]</i>
2022	LendOS founded. <i>[FACT]</i>
2022 (Sep 8)	Second Circuit reverses : Revlon lenders must return ~\$500m; discharge-for-value applies only where the debt is presently due. <i>[FACT]</i>
2022 (Sep)	AccessFintech Series C, \$60m. <i>[FACT]</i>
2022 (Oct 12)	Rehearing en banc denied. <i>[FACT]</i>
2022 (Dec 1)	The DLT venture files for Chapter 11 — having won its litigation. <i>[FACT]</i>
2022 (Dec)	Versana platform goes live. Revlon case settles and is dismissed. <i>[FACT]</i>
2023 (Sep)	Finastra closes \$5.3bn — a \$4.8bn unitranche plus a \$500m revolver at SOFR+7.25% at 98 OID — the largest US private-credit transaction to date. <i>[FACT]</i>
2023 (Feb)	Octaura completes the market's first fully electronic syndicated loan trades — a live-trade milestone, ahead of platform launch. <i>[FACT]</i>
2023 (Apr)	Octaura's syndicated loan trading platform officially launches . <i>[FACT]</i>
2023 (Jul)	LSTA folds distressed provisions into the par standard terms. <i>[FACT]</i>
2023 (May)	LSTA announces the Loan Identifier Cross-Reference — a translation service between competing identifier schemes. <i>[FACT]</i>
2023 (Oct)	Versana passes 1,500 facilities and ~\$900bn of commitments. <i>[FACT]</i>
2023	Average par settlement ~19.5 days. <i>[FACT]</i>

YEAR	EVENT
2024 (Jan 2)	KKR completes 100% ownership of Global Atlantic. <i>[FACT]</i>
2024 (Apr)	Citigroup credits an account with \$81 trillion instead of \$280; caught ~90 minutes later. <i>[FACT]</i>
2024 (Jun)	JPMorgan reports a straight-through-processing milestone on Versana. <i>[FACT]</i>
2024 (Sep)	Citi and Apollo announce a \$25bn direct-lending programme. <i>[FACT]</i>
2024 (Nov)	Finastra SFTP credential-based data breach — not ransomware; ~400GB claimed by the actor. <i>[FACT]</i>
2024 (Dec)	BlackRock agrees to acquire HPS for ~\$12bn. <i>[FACT]</i>
2024	US institutional loan issuance a record ~\$1.36tn, including a record ~\$760bn of repricings; US CLO issuance a record ~\$202bn; global syndicated lending \$5.9tn, +32%. CLOs buy 61% of new-issue leveraged loans. <i>[FACT]</i>
2025 (Jan)	Versana's first buy-side subscriber; reconciliation module launched. <i>[FACT]</i>
2025 (Feb)	JPMorgan increases a direct-lending commitment to \$50bn. <i>[FACT]</i>
2025 (Mar)	A data and research group acquires an AI-native private-credit monitoring platform positioned on source-cell traceability. <i>[FACT]</i>
2025 (Apr 2)	Finastra's \$5bn+ refinancing attempt collapses. <i>[FACT]</i>
2025 (May 19)	Finastra agrees to sell its Treasury and Capital Markets division. <i>[FACT]</i>
2025 (Jul 1)	BlackRock closes the HPS acquisition. <i>[FACT]</i>
2025 (Jul)	Finastra's refinancing revived and completed at ~\$3.6bn equivalent, oversubscribed. <i>[FACT]</i>
2025 (Sep)	First Brands files with liabilities of \$10–50bn (petition range; over \$9bn per the DOJ indictment). Tricolor collapses with \$900m+. Double-pledged collateral and fabricated receivables. Note: liabilities and collapse size are not losses — see Chapter 59.1. <i>[FACT]</i>
2025 (Jun 9)	Octaura announces the closing of a \$46.5m funding round. <i>[FACT]</i>
2025 (Sep 16)	Octaura CLO trading platform launches. <i>[FACT]</i>

YEAR	EVENT
2025 (Sep)	LendOS announces a Series A led by a major alternative manager's innovation arm; size undisclosed. <i>[FACT]</i>
2025 (Nov 24)	AccessFintech announces a strategic partnership with a major buy-side portfolio platform. <i>[FACT]</i>
2025 (Dec 5)	OCC and FDIC rescind the 2013 leveraged lending guidance , citing that it pushed lending into private credit. The Federal Reserve has not joined. <i>[FACT]</i>
2025	Secondary loan trading a record ~\$971bn. US CLO issuance ~\$209bn. Tricolor founder charged. <i>[FACT]</i>
2026 (Jan)	First Brands principals charged. Octaura at ~7.7% of secondary volume. <i>[FACT]</i>
2026 (Feb)	A major manager gates a retail debt fund. A debt-capital-management platform is acquired. An independent commercial-lending platform is acquired by a private-equity firm. <i>[FACT]</i>
2026 (Mar)	A flagship corporate lending fund imposes a 5% withdrawal limit. <i>[FACT]</i>
2026 (Q1)	Trailing-twelve-month secondary volume passes \$1tn for the first time. <i>[FACT]</i>
2026 (Apr)	Record ~6.0% US private-credit default rate. A major manager caps redemptions at 5%. <i>[FACT]</i>
2026 (Apr 30)	Versana closes \$43m led by BNP Paribas, adding Apollo; >\$125m total raised; explicit buy-side and private-credit push. <i>[FACT]</i>
2026 (May 6)	FSB publishes its Report on Vulnerabilities in Private Credit , naming valuation opacity and data gaps. <i>[FACT]</i>
2026 (May)	LSTA introduces late-payment fees for loan trade transactions. Citi and HPS agree a €15bn European direct-lending collaboration. <i>[FACT]</i>
2026 (Aug)	Vista explores a sale of Allvue at up to ~\$3bn. A leak-site campaign names Finastra; scope provisional. <i>[FACT/UNKNOWN]</i>
2026 (Sep)	2026 finance and banking layoffs reported at ~74,855 across nineteen firms, concentrated in back and middle office. <i>[FACT — INDICATIVE]</i>

Participant register

Organised by the problem each participant solves. Ownership as reported.

Entries the research could not verify are marked.

Agent and borrower servicing cores

PARTICIPANT	POSITION	OWNERSHIP
Finastra Loan IQ	Dominant syndicated and commercial loan servicing core; system of record for agents	Finastra (Vista Equity Partners)
Loan IQ Nexus	Interface and interoperability layer around Loan IQ; extended to private credit	Finastra
FIS Commercial Loan Servicing (ACBS)	Enterprise commercial and syndicated servicing; the market's other core	FIS (public)
FIS SyndTrak	Loan syndication and deal management	FIS
AFSVision	Real-time commercial lending platform	OceanSound Partners (announced Feb 2026)
nCino, Temenos, Mambu, Oracle FLEXCUBE, Fiserv	Origination and core-banking platforms serving adjacent segments	Various
LendOS	Cloud lending lifecycle platform for private credit; founded 2022	Private; Series A led by Blackstone Innovations Investments

Buy-side, CLO and private-credit portfolio systems

PARTICIPANT	POSITION	OWNERSHIP
S&P Global WSO	Loan and CLO portfolio administration; software and managed service	S&P Global
Allvue Systems	Front-to-back credit and CLO platform; ~50 of top 100 CLO managers	Vista Equity Partners (sale process, up to ~\$3bn)
Solvas	CLO accounting, compliance, priority of payments	Alter Domus (acquired from Deloitte)
Broadridge Sentry PM	Private-debt and CLO front-to-back	Broadridge (public)
Clearwater Analytics	Investment accounting, reconciliation, reporting	Public
SS&C	Fund administration plus loan and alternative-asset servicing technology	Public
"Everest"	Named in secondary sources; could not be verified as a distinct product	UNVERIFIED

Data, pricing, identifiers, analytics

PARTICIPANT	POSITION
S&P Global (loan pricing, LoanX/LXID, reference data)	Dominant loan data and identifier franchise; LXID extended to private credit June 2025
PitchBook / LCD	Leveraged-loan news, data and indices; divested from S&P Global to Morningstar in 2022
Bloomberg	Syndicated loans data, evaluated pricing, an open identifier standard, a leveraged loan index
Moody's, S&P Ratings, Fitch, KBRA, Morningstar DBRS	Ratings
CUSIP Global Services	CUSIP and the CUSIP Entity Identifier, developed with the LSTA and Versana
Kaneraï, KopenTech, 73 Strings	Structured-credit analytics, CLO electronic markets, valuation and extraction

Reconciliation and data-collaboration networks

PARTICIPANT	POSITION	OWNERSHIP
AccessFintech (Synergy)	Multi-party coexistence data network; dedicated loan data management; >250 institutions	Bank- and VC-backed
Versana	Bank-founded agent-data platform; real-time API capture; >\$4.1tn facilities	Bank consortium plus Apollo
DTCC Loan/SERV	Industry-owned reconciliation and messaging utility	DTCC
Duco, Gresham, SmartStream	General-purpose reconciliation platforms	Various

Settlement and post-trade

PARTICIPANT	POSITION	
ClearPar	Dominant electronic loan trade settlement; "substantially all" LSTA leveraged loans per a securities filing	S&P Global
Virtus	Competing settlement platform, founded 2005	Private
DTCC	Post-trade utility	

Electronic trading

PARTICIPANT	POSITION
Octaura	Bank-backed loan and CLO e-trading; RFQ and list protocols; ~7.7% of secondary volume (Jan 2026)
KopenTech	CLO auctions and peer-to-peer trading
MarketAxess, Tradeweb, Trumid	Adjacent corporate-credit venues (context)

AI, contract extraction and private-credit operations

PARTICIPANT	POSITION	STATUS
Hebbia, Oculus, Kira/Litera, Canoe, Cardo AI, Harvey, Luminance, AlphaSense	Document and contract extraction	Crowded; consolidating
Lumonic	AI-native private-credit monitoring with source-cell traceability	Acquired by Morningstar
73 Strings	Valuation-first AI extraction	\$55m Series B, Feb 2025
Setpoint	Capital-markets operating system; collateral verification, borrowing base, compliance	>\$50bn financed assets/yr
Finley	Debt-capital management	Acquired Feb 2026
Cascade, Casca	Asset-based finance ops; AI-native origination	Private
Evisort	Contract intelligence	Acquired by Workday

Administration, agency, trustee and BPO

Alter Domus (self-described largest loan agency provider, \$1.4tn administered) · IQ-EQ · SRS Acquiom · SS&C · Genpact · WNS · Wipro · Delta Capita · Wilmington Trust · U.S. Bank · BNY · Computershare · CSC · Kroll · S&P WSO Services · FIS Virtus

Consultancies and implementation

Accenture · Luxoft/DXC · Capgemini · TCS · Wipro · Cognizant · Infosys · Atyeti · Delta Capita · LS2 Partners · Adventure Consultancy Solutions

Trade bodies and standards

LSTA (1995) · LMA (1996) · APLMA (1998) · ARRC · FpML (ISDA) · CUSIP Global Services

Blockchain and DLT, including failed

Finastra Fusion LenderComm (Corda; status unknown post-2019) · Symbiont/Synaps (**bankrupt 2022**) · R3 · HQLAx (adjacent; migrating off Corda) · Provenance (loan-syndication relevance unverified)

Claims register

Every load-bearing claim in this paper, with its evidence class. A reader who disagrees with the paper's conclusions should locate the disagreement here.

FACTS the argument depends on

#	CLAIM	WHERE
F1	The agent maintains the register and it is authoritative on ownership	Ch. 1, 9, 68
F2	Loans are contracts, not securities; not reported to TRACE; no central depository	Ch. 1, 21
F3	Settlement standards are T+7 par / T+20 distressed; delayed compensation applies beyond	Ch. 22
F4	Average par settlement ~19.5 days (2023); 23% within T+7 (2016)	Ch. 22, 39
F5	Commercially significant credit-agreement terms remain bespoke; model forms rarely used as a starting draft	Ch. 36
F6	>7 million agent notices received a year by a single hub, >4 million redistributed (2010)	Ch. 12, 18
F7	CLOs bought 61% of new-issue leveraged loans in 2024 and own ~64%	Ch. 4, 10
F8	Citi/Revlon: ~\$893.9m wired 11 Aug 2020; SDNY ruling Feb 2021; Second Circuit reversal Sept 2022; settled Dec 2022	Ch. 8
F8a	OCC \$400m civil money penalty, Oct 2020, for enterprise-wide risk-management, data-governance and internal-control deficiencies — a separate action from the Revlon error, which is widely cited as emblematic of those deficiencies but is not established as its cause	Ch. 8, 76
F9	Citi \$81tn near-miss, April 2024	Ch. 8
F10	First Brands filed Sept 2025 with \$10–50bn of liabilities (petition range); Tricolor collapsed at \$900m+ — double-pledged collateral and fabricated receivables	Ch. 59
F10a	Documented lender figures are of four distinct kinds: gross exposure (~\$715m at one institution, >\$500m at another), charge-off (\$170m at a third), and realized hit (~\$10m after a prior \$30m at the first). No aggregate loss figure is supported and none is used in this paper.	Ch. 59.1
F11	FSB report 6 May 2026 names valuation opacity and data gaps	Ch. 56
F12	SEC 2026 exam priorities target private-credit valuation and liquidity	Ch. 56
F13	OCC/FDIC rescinded the 2013 leveraged lending guidance, 5 Dec 2025	Ch. 5
F14	Versana raised \$43m in April 2026 led by BNP Paribas, adding Apollo, with a buy-side and private-credit push	Ch. 46, 78
F15	AccessFintech: >250 institutions, >1bn transactions/month, 2021 Finastra integration, Nov 2025 buy-side partnership	Ch. 12, 46, 77
F16	The DLT venture won its litigation (\$53m settlement) and filed Chapter 11 on 1 Dec 2022	Ch. 45, 52
F17	FpML for loans published and not adopted; a matching product "has yet to win any users"	Ch. 38
F18	LSTA's identifier answer is a cross-reference, not a single identifier	Ch. 37

#	CLAIM	WHERE
F19	ABA Core Platforms Survey (fieldwork late 2024, released February 2025): 53% satisfied, ~35% dissatisfied, 69% likely to stay, ~19% expect to switch, satisfaction worst at renewal – core banking platforms, not loan servicing	Ch. 32, 77.9
F20	No public case of a bank migrating off Loan IQ; ACBS→Loan IQ flow is one-directional	Ch. 27, 31
F21	Finastra adjusted leverage ~16×; April 2025 refinancing collapsed; July 2025 revived and oversubscribed	Ch. 33
F22	Finastra is keeping Loan IQ and aiming it at private credit	Ch. 33, 60
F23	Vista owns both Finastra and Allvue	Ch. 33
F24	S&P Global holds ClearPar and WSO; antitrust remedies addressed indices and PRAs, not settlement or administration	Ch. 12
F25	ICSE 2025 study: LLM legacy-code comments "hallucination-free... useful"; and "no automated metrics strongly correlate with comment quality"; human evaluation necessary	Ch. 62, 63
F26	Strangler fig / event interception / ACL / parallel run are named, published patterns with documented risk profiles	Ch. 75
F27	Pattern guidance states the approach is unsuitable where requests cannot be intercepted or legacy source cannot be modified	Ch. 74, 77
F28	GLBA-type frameworks make the institution responsible for third-party safeguarding of NPI	Ch. 74
F29	Regulatory trajectory runs from screen scraping toward sanctioned APIs (PSD2 2018; US open-banking rule Oct–Nov 2024)	Ch. 74
F30	Protocol reverse engineering: network-only inference is inferior; encryption largely defeats it	Ch. 74
F31	Berg/Saunders/Steffen (2016): >80% of loans carry fees; fees can exceed interest	Ch. 14
F32	Mi & Han (2020): +4.15bp ≈ \$152k per unit of lead-arranger concentration; reverses for non-bank leads	Ch. 14
F33	2024 leveraged loans: BofA #1, JPM #2; top four ~30% of volume	Ch. 14
F34	Corporate bond electronic penetration ~43% in 2024 vs single digits in loans	Ch. 21
F35	Post-trade transparency reduced corporate bond bid-ask spreads substantially	Ch. 21, 53
F36	Only 17.7% of one trade body's members using or considering blockchain (Nov 2020) vs 60.4% for electronic platforms	Ch. 45
F37	LSTA's own position: "changing credit agreement provisions is necessary to shorten settlement times"	Ch. 22, 39, 86
F38	The agent's records are conclusive absent manifest error ; the post-Revlon erroneous-payment regime exists because agents make correctable mistakes	Ch. 68

#	CLAIM	WHERE
F39	Public security-interest filing regimes exist and give notice of competing claims; the 2025–26 frauds occurred with those regimes in place	Ch. 72, 73.1
F41	UCC Article 9 is a notice-filing system: what is filed is a limited record indicating that a person <i>may</i> have a security interest, and the filing office's role is ministerial with no substantive review	Ch. 73.1; App. F §B
F42	Filing, attachment, perfection and priority are distinct. § 9-502(d) permits filing before attachment; § 9-203(b) requires value, debtor's rights and an authenticated agreement; § 9-308(a) makes perfection attachment plus the perfection step; § 9-322(a)(1) ranks by first to file or perfect	Ch. 73.2; App. F §D
F43	The supergeneric asymmetry: § 9-504(2) permits an "all assets" indication in a financing statement; § 9-108(c) does not permit it in a security agreement	Ch. 73.2; App. F §G
F44	§ 9-204(a) permits after-acquired collateral, so a filing covers invoices that do not yet exist; § 9-210 information rights run to the debtor only	Ch. 73.2; App. F §G
F45	§ 9-318(a): a debtor that has sold an account retains no interest in it, so a later security interest cannot attach; and a financing statement does not disclose whether a transaction was a sale or a loan	App. F §J
F46	State title and electronic lien-and-title systems maintain authoritative per-vehicle lien records; no equivalent per-invoice record exists for trade receivables	Ch. 73.3, 73.4; App. F §F
F47	UCC Article 12 creates controllable electronic records and control-based perfection with super-priority, enacted in a majority of but not all states; it creates no registry and verifies no underlying obligation, and ordinary trade receivables are not controllable accounts absent deliberate tokenisation and the account debtor's undertaking	App. F §N
F48	Tricolor: prosecutors estimate ~\$2.2bn pledged against ~\$1.4bn of actual collateral, and Bloomberg reported ~29,000 of ~70,000 loans sharing identical vehicle identification numbers — allegations and reporting, not adjudications	Ch. 59.1, 73.3; App. F §Q
F49	Petters: \$3.65bn scheme, ~\$2bn of fabricated receivables, conviction and a fifty-year sentence — the strongest evidentiary class of the three cases	Ch. 73.3; App. F §Q
F40	Global CLOs outstanding ≈ \$1.2–1.5tn; US ≈ \$1.5tn; Europe ≈ €300bn — separately reported, on different measures and dates, and not additive or directly comparable	Ch. 4

VENDOR CLAIMS the paper reports but does not rely on

#	CLAIM
V1	Loan IQ: 21 of top 25 lenders, ~70% of global syndicated volume, 9 of top 10 agent banks, ~\$7.5tn facilities
V2	ACBS: "over 30% of syndicated loans worldwide"
V3	Post-close administration = 70–80% of the agent's operational burden
V4	McKinsey: 20–70% attrition; 1–2 years to productivity; 20–30% automation savings; ~85% manual
V5	Versana facility coverage >\$4.1tn
V6	Octaura share and volume figures (self-reported)
V7	AccessFintech scale metrics
V8	Allvue: \$7tn tracked; 50 of top 100 CLO managers
V9	Alter Domus: largest loan agency provider; \$1.4tn administered
V10	Private-credit AUM projections to \$5tn

INFERENCES that carry the argument

#	INFERENCE	DEPENDS ON
I1	The five-category complexity decomposition, and that AI changes only Category 2	Ch. 26, 67
I2	Standardization fails when cost falls on the agent and benefit on the lenders	Ch. 41
I3	The moat is Layers 2–5 (books and records, integration, muscle memory, career risk), untouched by product quality	Ch. 30
I4	The knowledge cliff makes migration harder, not easier	Ch. 66
I5	DLT answered consensus; the market's problem is provenance	Ch. 45, 71
I6	The market has no missing golden source; authority is distributed by design	Ch. 68, 70
I7	The rows of the authority table where no authority exists are the rows where the market's recent losses occurred — with the collateral row splitting into rank (authoritative), existence (authoritative and unconsulted) and uniqueness (genuinely unheld)	Ch. 69, 73.9, 81
I8	Verification produces information and evidence, not authority	Ch. 72
I9	The consortium response is triggered by success	Ch. 50
I10	The coexistence position is occupied and contesting it head-on is an uphill network-effects fight	Ch. 77
I11	The observe-and-report version prices as a reporting tool; the valuable version does not clear change management	Ch. 77
I12	Private credit is converging on the syndicated market's institutional forms, deliberately	Ch. 60, 78
I13	The industry's realistic authority model is many authoritative sources + governed synchronization + provenance; a single golden source is a category error	Ch. 70, 89
I14	Standardization splits into six layers with different prospects; the semantic layer is the highest-value one never seriously attempted, and the layer at which the Chapter 41 cost-benefit asymmetry is materially weaker — not absent : it avoids system replacement, franchise surrender, changes to negotiated economics, transfer of authority and universal-adoption dependence, but still requires governance, versioning, mapping, implementation, jurisdictional interpretation and change management	Ch. 84, 85, 91
I15	Translate where incumbency has created competing standards; standardize where nothing exists; allocate where the question is authority	Ch. 85
I16	Collateral integrity and private-credit valuation are different problems : the first has a verifiable external fact and a loss-prevention buyer, the second has neither and a structurally conflicted buyer	Ch. 69, 81.3, 94
I19	A legal system can allocate rights completely and correctly without establishing every underlying economic fact to which those rights relate; the precise form is that a facially valid public filing can exist without establishing that an enforceable security interest ever attached to the represented economic asset	Ch. 73.1, 73.2; App. F §D

#	INFERENCE	DEPENDS ON
I20	The residual collateral gap is primarily operational — authoritative mechanisms existed and were unused in every case examined — with one narrow structural component: uniqueness across facilities, which no single institution can establish alone	Ch. 73.3–73.6; App. F §§
I21	A shared uniqueness facility is non-linear in adoption: partial coverage yields a warning signal, only complete coverage of the relevant market yields a uniqueness proof, and the filing system's own history suggests universal coverage follows legal compulsion rather than voluntary agreement	Ch. 73.8; App. F §T
I22	Document provenance is not economic-asset existence — a hash of a fabricated invoice is still evidence of a fabricated invoice	Ch. 73.7; App. F §O
I23	Machine reading materially shortens the collateral pipeline — reading, extraction, obligor matching, contradiction detection, confirmation preparation, evidence assembly — and its output is a reason to doubt , not an establishment of fact: AI can detect a reason to doubt; it cannot manufacture authoritative evidence that does not exist	Ch. 63, 65, 88; App. F §P
I24	The incumbent's defensibility is an institutional ecosystem rather than a product — software, configuration, internal knowledge, external expertise, implementation firms, procedures, contractors, testing infrastructure and career capital — and a liquid external expertise market exists for the incumbent platforms and not for any challenger	Ch. 13, 30.1
I17	Four non-regulatory mechanisms exist to compensate the agent — fees, ownership, commercial incentive, contractual reform — of which ownership is proven and fees are unexplored	Ch. 86
I18	Steps 1–3 of the institutional sequence (interpretation, provenance, independent verification) require no other party's cooperation and are available now	Ch. 91

HYPOTHESES stated to be tested

H1–H6 (Ch. 75.4) · T1–T7 (Ch. 79) · the four predictions of Ch. 41.3 · the three residual uncertainties of Ch. 81.4 · the counterfactual that a verification system would have caught the 2025–26 frauds (Ch. 59.2) · that a second reading of the ACBS evidence implies both cores are simply immovable (Ch. 29)

UNKNOWNNS that materially limit the conclusions

#	UNKNOWN
U1	Any independent measurement of Loan IQ's market share
U2	The cost of a loan-servicing platform migration
U3	Whether any institution has granted a third party observation access to loan servicing traffic
U4	Loan IQ pricing, contract structure, release cadence, current architecture
U5	Whether Fusion LenderComm still operates
U6	Any published loan-operations error or break rate
U7	Median age or age distribution of the loan-operations workforce
U8	Current (2025–26) mean and median settlement days
U9	The par/distressed mix of 2025 secondary volume
U10	The aggregate cost of the market's duplicated processing
U11	Whether the ABA core-platform satisfaction curve holds for loan servicing specifically
U12	The scope of the August 2026 Finastra leak and whether loan data is implicated
U13	Any documented mechanism by which a shadow system acquires contractual authority
U14	Per-lender KYC cycle time and cost
U15	Whether institutions will procure collateral verification as a system, given that most of what failed was diligence they could already have performed — the post-failure reforms reported so far are contractual and procedural rather than infrastructural
U16	Whether the 2026 back-office reductions are caused by AI or merely attributed to it by the institutions announcing them
U17	Whether control under UCC Article 12 confers a uniqueness guarantee or only a priority rule
U18	Whether a cross-facility uniqueness obligation could be imposed by supervisory action without amending the Code
U19	The timetable on which the digital collateral infrastructure now under construction will reach warehouse lending, private credit and receivables, if it does

APPENDIX D

Figures this paper declines to use

Every figure below circulates in writing about this market. Each is excluded, with the reason.

FIGURE

WHY EXCLUDED

"847 material errors in 2025 causing \$142m in losses"	Traces to a single vendor page with characteristics of machine-generated content, sometimes misattributed to a trade association. Appears in no trade-body document.
"3–5% error rate on manual interest calculations"	Same source. No independent corroboration.
"200+ emails and 40–50 spreadsheet iterations per \$500m term loan B"	Vendor marketing; unverified; appears in at least two vendors' materials with no primary basis.
"75% fewer settlement errors from automation"	Vendor marketing.
"15,000–20,000 individual transactions annually"; "€185,000 discrepancy"; "Wells Fargo's operations team 60%..."	Same single dubious source.
"30% error rate on manual data entry / 50% excess headcount / 30–60 day cycles"	From a call-centre outsourcing blog, and about retail and mortgage processing rather than syndicated loans.
"\$2.5 trillion private credit market"	No independent source prints it. A defensible rounding of a range; an indefensible citation.
"Global syndicated loans market = \$682bn (2024) → \$783bn (2025)"	Market-research resellers with opaque methodology; appears to measure a narrow revenue segment; conflicts with regulator and trade-body data by an order of magnitude.
"\$1.8bn JPMorgan 2024 fees / \$29bn industry total / 588 lead arrangers"	Recurrs consistently in secondary reporting; could not be traced to a primary publisher. The bank publishes only its ranking.
"JPMorgan #1 for 26 consecutive years"	Reported in Chapter 14.2 as the bank's own claim citing a data provider, not declined outright. The #1 rank is well supported; the "26 years" phrasing is the bank's and is not independently verified.
"\$12.5 billion of losses" from First Brands and Tricolor	Adds First Brands' <i>liabilities</i> to Tricolor's <i>collapse size</i> and calls the sum a loss. Liabilities are what an estate owes; they are not lender losses, and documented lender exposures and charge-offs are two to three orders of magnitude smaller. This paper's own first draft made the error; Chapter 59.1 sets out the five distinct quantities.
"Twelve-figure" applied to ~\$12.5bn	Arithmetically wrong. \$12.5bn is eleven figures; twelve figures begins at \$100bn.
Global CLOs at \$1.4tn quoted alongside US CLOs at \$1.5tn	Both figures are separately reported and the pair is internally contradictory. They rest on different measures — notes outstanding, collateral under management, aggregate deal size — and different as-of dates. Chapter 4 now states the range and the reason.

FIGURE	WHY EXCLUDED
The OCC's \$400m penalty described as the penalty for the Revlon error	The enforcement action concerned enterprise-wide risk management, data governance and internal controls. The error is widely cited as emblematic of those deficiencies; the public record does not establish it as the cause of the fine.
"Four decades" of failed modernization	The market institutionalized between 1995 and 1998. It is a three-decade story, and this paper's own subtitle previously said otherwise.
A September 2023 date for Octaura's loan-platform launch	Circulates in secondary reporting and appeared in this paper's first draft. The correct sequence is: announced June 2022; first fully electronic syndicated loan trades February 2023; platform launched April 2023; CLO platform September 2025. The September 2023 date appears to conflate the loan platform with later CLO-platform coverage.
A June 2024 date for Octaura's \$46.5m funding round	Circulates in secondary sourcing and appeared in this paper's first draft. The round's closing was announced 9 June 2025.
"Nobody is authoritative about whether an asset was pledged twice"	Too broad as stated: public security-interest filing regimes exist precisely to give notice of competing claims and they rank them deterministically. The defensible claim is narrower still than the second draft made it, and is set out in Chapter 73: the filing system establishes rank rather than existence, and the fact no party holds is whether an asset sits in more than one lender's facility.
"CIPS has beaten SWIFT"	Conflates settlement value with messaging volume; the alternative rail remains ~80% dependent on the incumbent messaging network.
"Gold ~20% of reserves vs dollar ~46%"	A different denominator from the official reserve-currency series; not comparable to the ~58% figure.
"91% of CIOs..." credibility statistic	Surfaced in a search snippet; could not be confirmed on the source page.
"ClearPar settles 31% of trades"	Conflicts with the "substantially all LSTA leveraged loans" language in the same company's securities filing; context unverified.
"Markit RED" as a loan product	RED is a credit-derivatives reference-entity database. The loan analogue is LoanX ID. Misattributed.
"Loan IQ came from Midas"	Wrong. Loan IQ came from IQ Financial Systems, acquired from Deutsche Bank for \$40m.
"ACBS came to FIS via SunGard"	Wrong. It came via the ALLTEL/Fidelity acquisition completed 1 April 2003.
"Thoma Bravo owned FIS"	Unverified. The verifiable private-equity names are TPG and Thomas H. Lee (2005) and GTCR (via Metavante heritage and the 2023 Worldpay carve-out).
"The 2024 Finastra incident was ransomware"	Wrong. It was a credential-based SFTP data-exfiltration and extortion-sale incident; no malware, no encryption.
A named "Gary Wang" as a Loan IQ founder	No such person appears in any public Loan IQ source. Almost certainly a confusion with a verified, differently named consultancy managing partner.

FIGURE	WHY EXCLUDED
"Everest" as a buy-side CLO platform	Could not be verified as a distinct product.
"Mendel," "Sponsr," "Flow" as loan-tech ventures	Could not be substantiated.
A UK property lender's failure with a large collateral shortfall	Surfaced again in the collateral research, sourced to a vendor explainer and trade commentary rather than to a court, regulator or filing. Below the standard applied to the other cases; noted and not relied upon.
An aggregate loss figure across First Brands, Tricolor and a third case	Sums a reported liabilities figure, an alleged collateral shortfall and a reported shortfall at a fourth institution — three different quantities on three different measures, one of them not corroborated to this paper's standard. Appendix F states every figure separately with its class marked and adds none of them together.
"A perfectly perfected security interest existed over collateral that never existed"	Legally wrong, and it obscures the finding. A debtor has no rights in a nonexistent asset, so attachment fails under § 9-203(b)(2) and an unattached interest cannot be perfected. The correct statement is that a facially valid filing can exist without an enforceable security interest ever having attached.
"Article 9 failed"	Article 9 allocates legal rights and did so as designed throughout. The failures were of operational verification, not of the statutory scheme.
"No infrastructure exists to prevent double-pledging"	False as a general proposition. State title and electronic lien-and-title systems supply exactly that for vehicles, and were not queried. The claim holds for trade receivables and not for titled goods.
"UCC Article 12 creates an invoice-level registry"	It does not. Article 12 supplies a perfection-and-priority regime for controllable electronic records. It creates no registry, verifies no underlying obligation, and does not reach ordinary trade receivables.
2026 layoff counts	Directional moves are corroborated; the exact aggregate counts come from aggregator sites and are indicative rather than audited.

Open questions

The questions a fuller pass, or a reader with market access, should answer. They are grouped by what would be needed to answer them, because that determines who can close them and at what cost.

Legal counsel required

E1. Does control under UCC Article 12 confer a uniqueness guarantee, or only a priority rule? Whether the same controllable electronic record can be subject to multiple control arrangements is not resolved on the face of the statute. [U17]

E2. Could a cross-facility uniqueness-verification obligation be imposed by supervisory action — banking-agency guidance, or a securities rule for asset-backed issuance — without amending the Code? This determines whether the coordination problem of §73.8 has a route other than legislation. [U18]

E3. Does a lender's reliance on a borrower's compliance certificate without independent verification satisfy the good-faith standard, or can it be challenged as wilful blindness in a later avoidance action? [Appendix F §W]

E4. How is loss allocated, in a reorganisation following discovered double-pledging, among innocent secured creditors each holding a facially valid filing? [Appendix F §W]

E5. What quantum of "rights in the collateral" satisfies § 9-203(b)(2) for a pool of commingled after-acquired receivables? [Appendix F §W]

E6. Does the recent enforcement activity create a supervisory expectation of cross-facility collateral verification? If it does, the buyer in Chapter 94.1 acquires a compliance rather than a discretionary budget. [Appendix F §W]

Proprietary evidence required

E7. Has any institution ever granted a third party observation access to its loan servicing data flow, and on what terms? The single highest-value unknown in the paper. A positive answer materially improves the assessment of Path A. [U3]

E8. What does a loan-servicing platform migration actually cost, in money and elapsed time? The switching-cost argument is universally asserted and nowhere measured. [U2]

E9. Is there a published loan-operations error or break rate anywhere? The market's most-cited operational statistics are fabricated (Appendix D) and the real ones appear not to exist publicly. [U6]

E10. What is the aggregate cost of the market's duplicated processing? The market cannot measure its own central inefficiency. [U10]

Market validation required

E11. Will institutions procure collateral verification as a system, given that most of what failed was diligence they could already have performed? The reforms reported after the recent failures are expanded audit and confirmation rights — contractual and procedural, not infrastructural. This is now the question that determines the size of the opportunity in Chapter 94.1. [U15]

E12. Can collateral verification be delivered without servicing-core access? If yes, the best-evidenced proposition in the paper is also the most accessible. [T7]

E13. What does Versana's private-credit product actually do today, and which managers are live on it? This determines whether the window in Chapter 78 is closing or closed. [T6]

E14. Does the core-platform satisfaction-at-renewal curve hold for loan servicing? The paper's dissatisfaction argument rests on an analogy. [U11]

Research coverage gap

E15. What is Loan IQ's actual market share, from a source that is not its vendor? Every strategic document in this market rests on a marketing number. [U1]

E16. What are the current mean and median settlement days? The public record stops in 2023. [U8]

E17. What is extraction accuracy on real credit agreements, measured on economically operative terms rather than document-level metrics? [T2]

Other unknowns

E18. Are the 2026 back-office reductions caused by AI, or attributed to it? The institutions announcing them say AI; the public record does not distinguish cause from occasion, and the distinction matters for how durable the trend is. [U16]

E19. On what timetable, if any, will the digital collateral infrastructure now under construction reach warehouse lending, private credit and receivables? The direction is evidenced; the timetable is not. [U19]

Under-covered areas

No paper can cover the whole of institutional credit, and this one does not try. The following are named rather than filled, and the first is a deliberate scope decision rather than a gap:

- **Distressed credit — excluded by design.** The distressed segment operates under a different documentation regime, a different settlement window, its own par/distressed reclassification machinery, and a bankruptcy-court and claims-trading overlay that has no analogue in the par market. Chapter 24 describes the boundary and stops there. Compressing that material into this paper would do it badly; **a separate paper on distressed credit would be warranted if there is sufficient interest.**

- **Non-US market structure in depth.** European and Asian practice is described structurally but not analysed at the operational granularity of Part Three.
- **The trustee and corporate-trust layer.** Trustee market shares come largely from a competitor's ranking material.
- **Middle-market and bilateral lending.** The paper concentrates on broadly syndicated and private credit; the bilateral segment — roughly \$2.37tn of C&I loans at FDIC-insured banks in Q4 2024, and vastly more numerous by count — is a volume rather than a complexity problem and is treated only in outline. *[FACT]*
- **CLO trustee reporting and waterfall mechanics,** which are a substantial operational burden addressed only in passing.
- **Insurance statutory accounting,** which is a materially different set of requirements from fund accounting and is not developed.
- **The economics of the third-party administrator,** which Path B depends on and which this paper describes qualitatively.

Collateral integrity: legal authority, verification and the double-pledging problem

Legal-status note. This appendix analyzes secured-transactions law for the purpose of understanding the institutional operating problem; it is not a legal opinion. Statements labeled *[FACT]* reflect the cited statutory, judicial, regulatory, or other source. Conclusions about what those rules imply for market structure, verification, and technology are the authors' analysis and are labeled *[INFERENCE]* or *[HYPOTHESIS]* accordingly. Article 9 and related rules are subject to state enactment, asset-specific provisions, contractual terms, bankruptcy law, and other exceptions. Readers evaluating a particular transaction or legal position should rely on qualified counsel rather than this paper.

Chapter 73 carries the argument. This appendix carries the proof, the qualifications, the legal mechanics and the case detail, so that a reader who wants to check the chapter can do so and a reader who wants to disagree with it can find the place.

All statutory references are to the Uniform Commercial Code as promulgated by the American Law Institute and the Uniform Law Commission. Article 9 is enacted state by state and the enactments are not perfectly uniform; where this appendix says "the Code provides," it means the uniform text, and a transaction lawyer will check the enacting state. *[FACT]*

A. The question

Chapter 69 mapped the market's data elements to the party authoritative over each, and found three rows with no authority. Chapter 72 then made a claim about one of them — collateral — that was too broad in this paper's first draft and was flagged as an open question in the second.

The question this appendix answers is the narrow one: **how much of the collateral failure mode do existing security-interest filing regimes already address, and why did they not prevent the 2025–26 frauds?**

The answer, stated once and then supported: **the filing system addresses the allocation of legal rights, does so successfully, and was never designed to address the integrity of individual economic assets.** The residual gap is primarily operational, with one narrow structural component.

B. What Article 9 does

Article 9 governs the creation, perfection, priority and enforcement of security interests in personal property. Its filing system is a **notice** system.

The Official Comment is explicit. What is filed is not the security agreement but "a simple record providing a limited amount of information (financing statement). The notice itself indicates merely that a person **may** have a security interest in the collateral indicated. Further inquiry from the parties concerned will be necessary to disclose the complete state of affairs." *[FACT — § 9-502, OFFICIAL COMMENT 2]*

Two things a filing does establish:

WHAT IT ESTABLISHES	AUTHORITY
Constructive notice that the named secured party may claim an interest in the indicated collateral of the named debtor	§ 9-502, Official Comment 2
A priority date for first-to-file-or-perfect contests — priority dating from the earlier of first filing or first perfection	§ 9-322(a)(1) and Official Comment

Authorisation runs the other way and is worth stating for that reason: authenticating a security agreement itself constitutes the debtor's authorisation of a filing covering the described collateral. *[FACT — § 9-509(B), OFFICIAL COMMENT 4]* The security agreement supplies the authorisation; the filing does not evidence it, and an unauthorised filing sits in the public record looking exactly like an authorised one.

The filing office's function is ministerial. The grounds for refusing a record are enumerated and formal; none concerns the substance of the claim. The Official Comment states that a filing office may not reject a record even where the information it contains is facially implausible — the example given is an organisational identification number of obviously wrong form.

[FACT — § 9-516 AND OFFICIAL COMMENT]

C. What Article 9 does not establish

WHAT A FILING DOES NOT PROVE	AUTHORITY
That the collateral exists	§ 9-502 Official Comment 2; § 9-516 (no substantive review)
That the debtor has rights in the collateral	§ 9-203(b)(2) — a question of the debtor's property-law position, not of the filing
That the security interest attached	§ 9-203(b); § 9-502(d) permits filing <i>before</i> attachment
That the interest is perfected	§ 9-308(a) — perfection is attachment <i>plus</i> the perfection step
That the security agreement is enforceable	§ 9-203(b)(3)(A) — a matter of the agreement's own terms
That value was given, or how much	§ 9-502(a) — the form has no amount field
That the security interest was created at all	§ 9-203(b); <i>Raleigh Industries of America, Inc. v. Tassone</i> , 74 Cal. App. 3d 692 (1977): "filing is not necessary in order to create the security interest itself"

Article 9 also largely declines the ownership question: § 9-202 provides, subject to stated exceptions including certain sales of payment rights, that its rules apply whether title to collateral is in the secured party or the debtor. [FACT] Title is immaterial to the statute for most purposes, which is a deliberate design choice and a further reason the filing record cannot be read as an ownership register.

D. Filing, attachment, perfection, priority

These four are distinct and are the most frequently confused set of concepts in commentary about the recent failures.

CONCEPT	WHAT IT IS	WHAT SUPPLIES IT	VISIBLE IN THE PUBLIC RECORD?
Filing	The public act of recording a financing statement	The filing office, ministerially	Yes — and it may precede everything else (§ 9-502(d))
Attachment	The security interest becomes enforceable against the debtor	Value + debtor's rights in the collateral + an authenticated security agreement (§ 9-203(b))	No
Perfection	Attachment plus the applicable perfection step (filing, for accounts)	Both halves together (§ 9-308(a); § 9-310(a))	Only half of it
Priority	Rank among competing perfected claims	§ 9-322(a)(1), subject to exceptions such as purchase-money interests (§ 9-324) and buyers in ordinary course (§ 9-320)	The temporal ordering, yes; the exceptions, no

The operative consequence, stated precisely:

A facially valid public filing can exist without establishing that an enforceable security interest ever attached to the represented economic asset.

The mechanism for a fabricated receivable runs as follows. A financing statement is filed describing "accounts." A security agreement is authenticated. Value is advanced. But the debtor has no rights in a receivable that does not exist, so attachment fails under § 9-203(b)(2); and an interest that has not attached cannot be perfected. The filing remains on the public record and remains formally effective as a notice. There is no security interest behind it. *[INFERENCE — FROM THE STATUTORY TEXT]*

This is why the paper does **not** say that a perfectly perfected security interest existed over nonexistent collateral. That formulation is legally wrong and it obscures the finding, which is about the gap between a public record and an economic fact rather than about a defect in perfection doctrine.

E. The collateral fact matrix

Eleven facts a lender needs about an individual asset, against the channel that supplies each.

#	FACT	PUBLIC FILING	CONTRACT	BORROWER REPORTING	INDEPENDENT THIRD PARTY	AUTHORITATIVE PARTY	RESIDUAL GAP
1	Identity — this is the specific asset described	Class or type description only	§ 9-108 "reasonably identify"	Borrowing-base schedules	Field examination, sampling	Title registry for titled goods	No unique identifier exists for fungible goods, receivables or general intangibles
2	Existence — the asset exists	Nothing	Representations and warranties	Compliance certificates	Field audit, inspection, count	Account debtor (receivables); registry (titled goods)	For unregistered movables and intangibles, existence rests on borrower attestation plus periodic audit
3	Ownership — the debtor has rights sufficient to grant an interest	Nothing (§ 9-202 makes title immaterial)	Debtor's representation	Officer certificates	Lien search reveals competing filings	Registry for titled goods; intermediary for securities	No independent verification for general intangibles and after-acquired property
4	Uniqueness — pledged to one facility only	Records liens, does not cross-reference facilities	Negative pledge	Attestation	Cross-facility audit (rare)	None	The structural gap. No public or private mechanism at scale
5	Priority — the claim ranks where the lender thinks	§ 9-322(a)(1) filing date	Intercreditor agreements	n/a — a legal conclusion	Search report	The filing office, for temporal rank	Statutory exceptions and cross-jurisdictional gaps remain
6	Eligibility — qualifies under the advance criteria	Nothing (a private test)	Credit agreement defines it	Borrowing-base certificate	Field examiner tests against the contract	None; the standard is lender-defined	Depends on accurate borrower data on aging, dilution and quality

#	FACT	PUBLIC FILING	CONTRACT	BORROWER REPORTING	INDEPENDENT THIRD PARTY	AUTHORITATIVE PARTY	RESIDUAL GAP
7	Provenance — chain of title or creation history	Nothing	May be warranted	Records if required	Title chain for titled goods	Registry for titled goods and registered IP	No chain-of-title infrastructure for most personal property
8	Valuation — worth what is claimed	Nothing	Advance rates, appraisal requirements	Borrowing-base and financials	Appraisal, mark-to-market	Exchange price for traded securities	Inherently a judgment for private, illiquid or intangible assets — see §87.2 and Chapter 94.2
9	Monitoring — changes detected in time to act	Continuation every five years; no monitoring	Notification covenants	Periodic reporting	Periodic field exams	None	Event-driven and backward-looking; real-time only for electronically tracked assets
10	Borrower reporting — the representations are accurate	Nothing	Reporting covenants and certifications	The certificates themselves	Audit and examination testing	None; the borrower is the primary source	Structural asymmetry. Verification is periodic and sample-based
11	Fraud detection — misrepresentation caught before loss	None; the office is ministerial	Remedies, not detection	Analytics on borrower-supplied data	Surprise exams, third-party data matching	Enforcement, ex post	No preventive mechanism in the statute or in current market infrastructure

[FACT FOR THE STATUTORY COLUMNS; INFERENCE FOR THE RESIDUAL-GAP COLUMN]

Of the eleven, exactly one — priority — is delivered by public infrastructure. Three — existence, eligibility and borrower reporting — are addressable by an individual institution's diligence, periodically and on a sample. Five are partially addressed and depend on asset class. One — uniqueness — is not addressed by anything currently operating at scale for receivables. And one — preventive fraud detection across the whole

collateral base in real time — the research classes as not independently establishable at all with what exists today, which is a stronger statement than "unsolved" and is the reason row 11 of the table above says what it says. *[INFERENCE]*

F. Asset classes differ, and the difference is decisive

ASSET CLASS	IS THERE AN AUTHORITATIVE PER-ASSET RECORD?	CONSEQUENCE
Titled goods (motor vehicles)	Yes — state title and electronic lien-and-title systems record the noted lienholder per vehicle	The uniqueness question is answerable today. Failure to answer it is operational
Securities held through the clearing system	Largely — intermediary and depository records establish holdings	Ownership is recorded; tracking of collateral reuse across a pledge chain is a separate and less settled problem
Registered intellectual property	Partially — registration records exist; security interests are recorded inconsistently	Partial
Inventory and equipment	No — no per-unit registry	Physical inspection is the substitute
Trade receivables	No — no per-invoice record anywhere	The account debtor is authoritative on the underlying obligation; nobody is authoritative on how many times it has been financed

This table is the single most important qualification in the appendix. A claim that "no infrastructure exists to prevent double-pledging" is false as a general proposition. For a substantial class of assets it exists, is authoritative, and is operational. *[INFERENCE]*

G. Receivables and the granularity problem

Four features of the statute, taken together, explain why the filing system cannot be made to answer invoice-level questions.

One: notice-filing is designed for class descriptions. A single financing statement is intended to cover transactions under agreements not yet made and collateral not yet acquired. *[FACT — § 9-502, OFFICIAL COMMENT 2]*

Two: the supergeneric asymmetry. § 9-504(2) permits a financing statement to indicate that it covers "all assets or all personal property." § 9-108(c) provides that the same description in a *security agreement* does not reasonably identify the collateral. [FACT] The public layer is permitted to be less specific than the private one. The financing statement therefore establishes the outer boundary of what could be covered; the security agreement determines what is.

Three: after-acquired property. § 9-204(a) permits a security interest in after-acquired collateral, and the Official Comment confirms it is no less valid than an interest in collateral the debtor already holds. [FACT] Every future invoice falls automatically within the lien. It cannot be identified at the time of filing because it does not yet exist.

Four: information rights run to the debtor. § 9-210 gives a mechanism to request a statement of account and a list of collateral. Official Comment 3 states that a third party is not entitled to request it, on the reasoning that this would burden the secured party and that a third party can ask the debtor. [FACT] The party with the strongest interest in knowing what is already pledged has no statutory means of asking the party who knows.

What a searcher can therefore determine: that a party has filed against the debtor for accounts; the filing date; the secured party's identity. What a searcher cannot determine: which invoices are in a borrowing base; whether a specific invoice has been sold or pledged; the outstanding balance of the secured obligation; whether the interest attached; whether the receivables exist. [FACT]

The link between a specific invoice and a specific financing exists only in the borrowing-base certificate, the collateral schedule, and — where notification has been given — the § 9-406(a) notice to the account debtor. None of these is public. [FACT]

H. Duplicate financing: six cases

CASE	SITUATION	WHAT A LIEN SEARCH SHOWS	WHAT IT DOES NOT SHOW	WHERE THE FAILURE SITS
A	Two lenders, same broad collateral class, both filings valid	Both filings, with dates	Amount owed, which receivables, whether the debtor owns any	Nowhere. The junior lender's exposure is a priced credit decision
B	Debtor represents the same receivables as unencumbered to a second lender	Both filings	Whether the specific asset offered is already pledged	Debtor fraud. Priority resolves the contest; it does not create collateral. The recent auto and receivables cases sit here
C	Same receivable sold outright to a factor and pledged to a lender	Both filings — and the record does not distinguish a sale from a pledge	Whether the debtor retained ownership	§ 9-318(a): a seller of an account retains no interest, so once the buyer has perfected the later interest has nothing to attach to (§ 9-203(b)(2)). Where the buyer has <i>not</i> perfected, the seller is treated as retaining rights and the contest reverts to ordinary priority. This is an ownership failure, not a priority failure
D	Fabricated receivable financed by several parties	Multiple formally valid filings	Anything about existence	Attachment fails as to every one of them. Priority over nothing is nothing
E	Same receivables placed in two securitization structures	Filings, but a true sale may move the debtor of record to the special-purpose entity, so a search against the originator may not surface the first transfer	Whether the same asset sits in two pools	Originator fraud; trustee, servicer and auditor verification obligations
F	Junior lender correctly understands its rank but misjudges the collateral cushion	Both filings and its own position	Senior balance, future-advance rights, current collateral value	Not a legal failure at all. An economic misjudgment, which courts do not relieve

[FACT FOR THE STATUTORY ANALYSIS; THE CASE-TO-CATEGORY ASSIGNMENTS ARE INFERENCE]

The taxonomy matters because it separates the cases Article 9 fully resolves (A, F), the cases where it resolves the legal contest but not the economic outcome (B, E), the case it resolves through the attachment requirement but invisibly to a searcher (C), and the case where it has nothing to work on (D).

I. Borrowing-base verification

The borrowing base is where the collateral question is actually operated, and it is entirely private.

Eligibility criteria — aging, cross-aging, concentration limits, dilution thresholds, disputed accounts — are defined in the credit agreement. The Code does not regulate them. The borrower certifies compliance; the lender tests it. *[FACT]*

The verification channels, ordered by independence from the borrower:

EVIDENCE	INDEPENDENCE	COMMENT
Invoice	Low	Created by the borrower. Fabricated at scale in the recent cases
Purchase order or contract from borrower files	Low–moderate	Rises substantially if obtained from the counterparty instead
ERP data, aging, payment history	Low–moderate	Borrower-generated; cross-referencing patterns adds value
Shipping and delivery documentation	Moderate	Higher if sourced from the carrier
§ 9-406(a) notification and the account debtor's response	Moderate–high	A denial or a request for proof of assignment is itself a signal
Field examination	High	Tests sales, billing, write-offs, turnover, concentration; reconciles deposits to accounts
Lockbox and cash dominion	High	The borrower cannot fabricate bank-cleared receipts from third parties
Direct account-debtor confirmation	High	External confirmation is a recognised high-reliability evidence category in auditing standards

[FACT FOR THE CATEGORIES; THE INDEPENDENCE RANKING IS INFERENCE]

None of these is required by Article 9. The statute requires that the debtor have rights in the collateral for attachment to occur; it says nothing about how a lender should satisfy itself that this is so. *[FACT]* Verification is credit practice and supervisory expectation, not statutory obligation.

J. Factoring and receivable sales

Article 9 applies to outright sales of accounts as well as to security interests in them (§ 9-109(a)(3)), and uses the term "security interest" for both. *[FACT]* Two consequences follow that matter here.

First, a **financing statement does not disclose whether the underlying transaction is a sale or a loan**. A searcher sees a filing; the character of the transaction is in the documents. *[FACT]*

Second, § 9-318(a) provides that **a debtor that has sold an account retains no legal or equitable interest in it**. *[FACT]* Where the buyer has perfected, a later secured party therefore has nothing to attach to, because the debtor has no rights to convey; it may hold a perfect filing and no interest at all — the same result as with a fabricated receivable, arrived at by a different route. Where the buyer has not perfected, the Code treats the seller as retaining rights for this purpose and the contest reverts to ordinary priority, which is itself a trap for a purchaser who assumes the sale alone protects it. *[FACT]*

This is the mechanism by which a receivable that has been factored can be pledged again without any public record showing the conflict, and it is a materially different problem from ordinary lien priority.

K. The authority map for collateral

QUESTION	AUTHORITATIVE PARTY	AVAILABLE TO A LENDER?
Rank among competing secured parties	The filing system, by statute	Yes, publicly
Whether an invoice represents a genuine obligation	The account debtor	Yes, by direct confirmation
Whether a vehicle is already subject to a noted lien	The state title / electronic lien-and-title system	By query, where the lender participates in the state's programme or uses a title service; access to motor-vehicle records is restricted
Whether a receivable was previously sold	The purchaser and the seller's own records	Not to a third party; § 9-210 runs to the debtor only
Whether the same asset sits in another lender's facility	Nobody , for asset classes with no per-asset record; the title record answers it for a titled good	No
Whether a borrowing-base certificate is accurate	The borrower certifies; the field examiner tests	Partially, periodically, on a sample

[FACT FOR ROWS ONE TO THREE; INFERENCE FOR THE REMAINDER]

Read against Chapter 69, this table is the corrective. The authority is not uniformly absent. It is present and public for rank, present and reachable for existence, and absent only for uniqueness across facilities.

L. The case for existing diligence, at full strength

This is the strongest objection to the chapter and it is set out here without softening.

Every control needed to catch these frauds already existed and was well known. Post-mortem commentary from practitioners advances the further proposition that independent collateral verification was once ordinary market practice in asset-based lending and was competed away under margin pressure; the paper reports that as commentary and does not rely on it, because the sourcing available to this pass is legal-press analysis rather than a measured decline in verification practice. *[VENDOR CLAIM /*

UNKNOWN] Field examinations are prescribed in supervisory handbooks for asset-based lending portfolios. Unannounced audits, reconciliation of borrower data tapes against independent servicing records, cash dominion and lockbox arrangements, third-party custody of chattel paper, and direct confirmation with account debtors are all long-established, none of them novel, and all of them available to any lender that chose to use them.

The record supports the objection. In each of the three cases, warning signs were present and were not acted on: audits that surfaced discrepancies and were explained away, information requests routed through the people conducting the scheme, self-custody of the underlying paper permitted where third-party custody is available, and, in the auto case, no independent check of a per-vehicle title record of the kind that post-mortem analysis says would have exposed the duplication.

Verdict: substantially correct. Better single-institution diligence would have reduced or prevented the loss in every case examined. Any account of these events as purely structural overstates. *[INFERENCE]*

M. Where individual diligence stops

The objection is nonetheless incomplete, and the boundary is precise.

Diligence is a capability one institution applies to its own exposure. Uniqueness is a property of the relation between institutions. A lender that verifies its own borrowing base perfectly still cannot observe that the same receivable has been pledged into a facility at a different funder: it cannot see that facility, cannot compel its disclosure, and — for an asset class with no per-asset registry — no third party holds the union of the two. *[INFERENCE]*

Existence, eligibility and the accuracy of borrower reporting are all addressable by one institution acting alone at a cost it can choose to bear. Valuation is not on that list; it is a judgment rather than a fact and is treated separately (§87.2 and Chapter 94.2). Uniqueness across facilities, in

the asset classes that lack a per-asset record, is not addressable alone either. That is the whole of the structural claim, and the paper makes no wider one.

N. Does UCC Article 12 create an invoice-level registry?

No. The answer is unqualified and it is worth stating plainly because the contrary is frequently implied.

Article 12, adopted by the sponsoring bodies in 2022 and enacted in a majority of states, creates the category of **controllable electronic record** — a record in an electronic medium that can be subjected to control under § 12-105. *[FACT]* Control means the power to obtain substantially all the benefit of the record, the exclusive power to prevent others from doing so, and the exclusive power to transfer control. The paradigm assets are cryptocurrencies, tokens and digitally native instruments. *[FACT]*

The conforming Article 9 amendments add **controllable accounts** (§ 9-102(a)(27A)): an ordinary account that is evidenced by a controllable electronic record and where the account debtor has undertaken to pay whoever controls that record. *[FACT]* Perfection by control is available for these, and a secured party perfected by control takes priority over one perfected only by filing, regardless of temporal order — a departure from first-to-file-or-perfect. A qualifying purchaser takes free of prior property claims. *[FACT]*

What Article 12 does **not** do:

- It does not create, require or enable any registry of receivables, invoices or controllable electronic records — public, private, central or distributed. *[FACT]*
- It does not verify that an underlying commercial obligation exists. Control is a perfection concept; it is not an existence concept. *[FACT]*
- It does not establish real-time visibility into invoice status, payment or assignment. *[FACT]*

- It does not convert ordinary trade receivables into anything. An ordinary invoice is not a controllable electronic record; it becomes a controllable account only through deliberate tokenisation *and* the account debtor's undertaking. That is a structure the parties must build, not one that arises from ordinary trade. *[FACT]*

Control is also demonstrated bilaterally between transacting parties.

No central or public record of individual receivables is created or consulted. *[FACT]*

Two further limits belong in an honest statement. Enactment is state by state and is not complete, which affects nationwide enforceability during the transition. *[FACT]* And whether control by one party necessarily excludes control by another — that is, whether the control concept supplies a uniqueness *guarantee* rather than a priority rule — is not resolved on the face of the statute and is a question for counsel. *[UNKNOWN]*

O. Existing digital and registry infrastructure: what it actually does

CLASS OF SYSTEM	WHAT IT ESTABLISHES	WHAT IT DOES NOT
Government clearance e-invoicing (several jurisdictions operate mandatory platforms through which invoices pass for validation)	That an invoice was issued in valid form and reported to the tax authority; in at least one such platform, detection of duplicates submitted to it	Nothing about assignment, pledge or financing. These are tax-compliance systems
Trade-finance duplicate-financing registries (bank utilities that compare hashed document fingerprints)	That the same document has been presented by another participant	Nothing about non-participants, and nothing about whether the document is genuine
Receivables-exchange platforms requiring buyer confirmation before financing	That the named buyer confirms the obligation, and that it has not been financed twice <i>within the platform</i>	Nothing outside the platform. Coverage is a fraction of the addressable market
Supply-chain finance platforms operating against buyer-approved payables	That the buyer has approved the payable	Nothing about a separate pledge to a lender outside the programme
Electronic bills of lading	Title and transfer within a platform	Cross-platform transfer remains a work in progress; documents of title are not trade receivables
A tokenised-collateral platform under construction at a central clearing utility	Announced capability for centrally cleared securities collateral	Not live at the time of writing; not addressed to warehouse lending, private credit or receivables
Vendor duplicate-pledge detection products	Self-reported capability to fingerprint assets and alert funders to matches across facilities	Coverage limited to participating originators and funders. <i>[VENDOR CLAIM]</i>

[FACT FOR THE DESCRIPTIONS OF THE PUBLIC SYSTEMS; VENDOR CLAIM WHERE MARKED]

Two conclusions follow, and both belong in the paper rather than only here.

A hash of a fabricated invoice is still evidence of a fabricated invoice.

Establishing that a document exists, is unaltered and has been seen before is provenance, and provenance is useful. It is not existence of the economic asset, and the two are routinely conflated. *[INFERENCE]*

The gap is being addressed and is not closed. Infrastructure is under construction; none of it currently covers the asset classes and market segments where the recent losses occurred; and a claim that it will close

on a stated timetable is not supportable from the public record. [INFERENCE / UNKNOWN]

P. What AI can detect, and what it cannot establish

Every case in §Q turned on a fact that was knowable and was not known. That makes this the natural place to ask what machine reading changes, and the honest answer is narrower than the enthusiasm and larger than nothing. It is the argument of Chapters 63 and 65 applied to this domain. [INFERENCE]

The pipeline AI genuinely shortens. Six steps, each of which is today manual, sampled and periodic:

STEP	WHAT IT IS	WHAT CHANGES
Document reading	Ingesting invoices, schedules, purchase orders, shipping documents, borrowing-base certificates and the credit agreement's eligibility definitions	Was economically impossible at portfolio scale in 2022; is not now (Chapter 62)
Extraction	Structuring each document into fields — obligor, amount, date, terms, reference numbers	The unit cost falls far enough that <i>every</i> item can be read rather than a sample
Obligor matching	Resolving names and entities across documents, systems and facilities, including accented and ASCII variants	Entity resolution at scale is a machine problem, and doing it badly is itself a documented failure mode
Anomaly and contradiction detection	Finding items that disagree with each other, with the aging, with the payment history, or with the eligibility definitions the agreement actually states	This is the step where the value concentrates
Confirmation preparation	Assembling the population, the questions and the packets for direct confirmation with account debtors	Converts confirmation from an expensive sample into an affordable census
Evidence assembly	Retaining what was checked, against what, when, by what method, with what result	The provenance discipline of Chapter 71.1, applied to collateral

[INFERENCE, EXTENDING CHAPTERS 63 AND 65 TO THIS DOMAIN]

The step that matters most is the fourth, and it is worth being precise about what it produces. Cross-document contradiction is not a judgment about whether an invoice is genuine. It is the observation that two documents that should agree do not — an invoice with no matching shipping record, an

obligor whose payment history does not match its claimed balance, a receivable appearing in two schedules with different reference numbers, an eligibility test that the certificate reports as passed and the underlying data does not support. Each of those is a **reason to doubt**, and a reason to doubt is exactly what the recent failures lacked: not a conclusion, but a prompt to make the phone call that nobody made. *[INFERENCE]*

What AI cannot do here, and the boundary is hard.

It cannot establish that a document corresponds to a real transaction.

Only the counterparty to that transaction can, and a machine's contribution is to identify which counterparties to ask and to prepare the asking. A perfectly extracted fabricated invoice is a perfectly extracted fabricated invoice.

It cannot see into another lender's facility. The uniqueness gap of \$M is an access problem, not an inference problem. No amount of reading the documents one lender holds reveals what a different lender holds, and a model asked to infer it will produce a probability where the question demands a fact. *[INFERENCE]*

It cannot convert a probabilistic finding into an authoritative one. Chapter 65's boundary applies unchanged: the model proposes and something else proves. In this domain the prover is the account debtor, the title registry, the custodian or the bank statement — parties and records outside the model — and the model's output is admissible as a reason to consult them, not as a substitute for having done so. *[INFERENCE]*

And it cannot create evidence that does not exist. Where no authoritative record exists — the per-receivable case — no analytical technique produces one. That is a structural fact about the domain, and it is the second half of the finding this appendix exists to establish. *[INFERENCE]*

The boundary, stated once:

AI can detect a reason to doubt. It cannot manufacture authoritative evidence that does not exist.

Two consequences follow for anything built here, and they are the reason Chapter 94.1 is assessed as it is.

The first is that the achievable product is a **triage and evidence system**, not an oracle: it reads everything, flags what disagrees, prepares the confirmations, and keeps the record of what was checked. That is genuinely valuable, addresses the operational half of the gap identified in §S, and requires nobody else's cooperation. [INFERENCE]

The second is that its output is only as good as the confirmations it prompts. A system that flags anomalies into a queue that nobody works reproduces the failure it was bought to prevent — which is the same finding as §M, arriving from the technology side rather than the diligence side. [INFERENCE]

Q. The three cases, with a note on quantities

The paper's discipline on numbers (Chapter 59.1) applies here with particular force, because the figures circulating around these events measure at least six different things. Every figure below is labelled:

These quantity classes are **not** the paper's five evidence labels and are not substitutes for them; they describe what a number measures and how firmly it is established, and the ordinary labels continue to apply alongside them.

- **ALLEGED** — asserted in a civil complaint or an indictment, not adjudicated
- **CHARGED (criminal) / CHARGED (civil)** — the subject of criminal charges, or of a regulator's civil enforcement action; the distinction matters and is kept
- **FOUND** — established by conviction, judgment or admitted in a court-approved resolution

- **REPORTED** — reported by estate professionals, counterparties or the press, not tested
- **LOSS** — a realized loss or a charge-off recorded by a named institution

No figure in this appendix is added to any other by this paper. The three cases are not summed, and the paper uses no aggregate loss figure for them. Where a cell states a figure that is itself a difference or a total, the arithmetic was performed by the source and the cell says so.

First Brands Group

QUANTITY	CLASS	FIGURE
Liabilities at filing	REPORTED	\$10–50bn per the bankruptcy petition; over \$9bn per the DOJ indictment; ~\$11.6bn is an analyst total including off-balance-sheet financing
Fabricated or inflated receivables	REPORTED / ALLEGED	~\$2.3–2.7bn depending on the measure cited; the lower figure is described as off-balance-sheet factoring exposure and the higher as fabricated accounts receivable, and the paper does not treat them as interchangeable
Invoice inflation	REPORTED	Up to ten times face value in cited instances
Named counterparty exposures	REPORTED	Approximately \$1.75bn at one financing counterparty; ~\$763m at a second; ~\$172m at a third; a fourth undisclosed. These are exposures, not losses
Documented lender realized outcomes	LOSS	Chapter 59.1 carries the figures: a ~\$10m realized hit at one institution after a prior \$30m, having written the exposure to zero. The documented realized figures at named institutions are two to three orders of magnitude below the headline
Criminal proceedings	CHARGED (criminal)	Principals charged, January 2026, as recorded in Chapter 59.1

Failure categories: fabricated receivable; duplicate pledge; borrowing-base eligibility manipulation. **Why standard diligence did not catch it:** field examinations received aggregated schedules rather than original invoices; a 2023 audit finding was attributed to employee error; information requests were routed through the executives conducting the scheme. *[FACT — AS REPORTED]* **Authoritative source available:** yes — the account debtors, several of them large national businesses, could have confirmed or denied each invoice. Not used at scale. *[FACT — AS REPORTED]*

Tricolor Holdings

QUANTITY	CLASS	FIGURE
Collateral pledged against actual collateral	ALLEGED	~\$2.2bn pledged against ~\$1.4bn actual; the ~\$800m gap is the difference between them and is stated as such in the sources
Loans sharing identical vehicle identification numbers	ALLEGED	~29,000 of ~70,000 active loans
Single documented instance	ALLEGED	3,225 identical loans directed into two separate securitizations
Securities issued	REPORTED	~\$1.9bn through asset-backed offerings
Principal outstanding at bankruptcy	REPORTED	~\$945m+
Named lender figures	REPORTED / LOSS	~\$170m at one institution, recorded in Chapter 59.1 as a charge-off; at a second, reported ranges differ across sources at ~\$170–200m and ~\$178–200m and both are stated here rather than reconciled; ~\$50m and ~\$30m at two others. One source reports aggregate warehouse-lender reserves above \$500m — that aggregation is the source's, not this paper's. Reserves, charge-offs and exposures are different quantities and are not summed here
Proceedings	CHARGED (criminal) / CHARGED (civil)	Founder criminally charged, December 2025, as recorded in Chapter 59.1; securities-regulator civil charges against former executives, 2026 (SEC Press Release 2026-77)

Failure categories: duplicate pledge; borrowing-base eligibility manipulation through relabelling of delinquent loans. **Why standard diligence did not catch it:** loan-level data came from the borrower's own tapes; no lender performed cross-facility deduplication against independent data. *[FACT — AS REPORTED]* **Authoritative source available: yes, and this is the case's importance.** State title records, and electronic lien-and-title systems in the states that operate them, maintain per-vehicle lien records. The mechanism existed, was authoritative, was operational, and was not queried. That either it or a cross-facility comparison of vehicle

identification numbers would have exposed the duplication is the assessment of trade and legal-press analysis rather than a court finding.

[FACT — AS REPORTED]

Petters Group Worldwide

QUANTITY	CLASS	FIGURE
Total scheme	FOUND	~\$3.65bn
Fabricated receivables	REPORTED	~\$2bn — secondary reporting; only the scheme size and the sentence attach to the conviction
Sentence	FOUND	Fifty years, 2010
Fund manager	CHARGED (civil)	The manager of a fund that purchased the notes, charged by the securities regulator for facilitating the fraud by failing to perform basic verification

Failure categories: fabricated receivable; existence fraud — the collateral did not exist at all. **Why standard diligence did not catch it:** it was not performed. Purchase orders purportedly issued by two named national retailers were accepted without confirmation. **Authoritative source available:** yes — the named retailers. A written confirmation request would have exposed the scheme.

On a fourth case. A reported failure at a United Kingdom property lender, with a collateral shortfall in the hundreds of millions of pounds, appears in the collateral research and in trade commentary. The paper does not rely on it: its sourcing is a vendor explainer and trade press rather than a court, a regulator or a filing, which is below the standard the paper applies elsewhere. It is recorded in Appendix D among the figures declined. [UNKNOWN]

R. Comparative failure-mode matrix

	FIRST BRANDS	TRICOLOR	PETTERS
Asset class	Trade receivables	Auto loans and the vehicles securing them	Purported consumer-goods receivables
Existence failure	Yes	No — the loans existed	Yes — total
Duplicate pledge	Yes	Yes	No
Eligibility manipulation	Yes	Yes	n/a
Would a lien search have detected it?	Partially — entity-level overlap visible, invoice-level duplication invisible	Partially — blanket liens visible, loan-level duplication invisible	No — a filing says nothing about existence
Would standard diligence have detected it?	Partially. Obligor confirmation would have exposed the fabrication; it would not have exposed a genuine invoice sold twice	Partially. A title query or cross-facility comparison of vehicle identification numbers would have exposed the duplication; neither addresses the relabelling of delinquent loans as current	Yes — a confirmation request would have exposed it
Was an authoritative source available?	Yes — the account debtors	Yes — state title records	Yes — the named retailers
Was it used?	No	No	No
Residual gap after perfect single-lender diligence	Uniqueness across facilities	Minimal — the verification mechanisms existed	None

[INFERENCE, FROM THE CASE RECORDS]

The last row is the finding. In two of the three cases, competent diligence by a single institution would have closed the gap, because the authoritative record or the authoritative counterparty was reachable. In the third — the receivables case — it would have closed most of it and left one question standing, and that question is the only structural claim this paper makes.

S. Legal, operational or infrastructural?

LAYER	VERDICT
Legal	Not the gap. Article 9 allocates rights as designed. The recent failures did not turn on a defect in attachment, perfection or priority doctrine
Operational	The primary gap. Mechanisms existed — entity-level lien searches, per-vehicle title queries, account-debtor confirmation, third-party custody, cash dominion — and were not used
Infrastructural	A narrow residual. For receivables there is no per-asset registry, no cross-lender deduplication service, and no independent continuous borrowing-base verification. These do not exist rather than existing and being ignored

The distinction between "exists and was not used" and "does not exist" is the analytical core of this appendix, and it is the distinction most commentary on these events omits. *[INFERENCE]*

T. The coordination test

Applied to the one mechanism that requires other institutions: a shared uniqueness facility.

DIMENSION	ASSESSMENT
Who must participate	Every lender extending credit against the same asset class to the same borrower family — including non-bank funders, factors and cross-border lenders
Effect of one non-participant	The non-participant becomes the channel. One defector converts a proof into a warning
Is universal adoption required	For a uniqueness proof, yes. A duplicate is visible only if <i>both</i> pledges are inside the network, so coverage compounds: at ninety-five per cent participation roughly nine in ten duplicate pairs are visible, and every pair with one leg outside is invisible
Network dynamics	Value is non-linear in adoption, with a threshold below which the facility is advisory only
Precedent	The filing system itself. Notice filing achieved universal coverage because filing was made a legal prerequisite to priority, not because participants agreed it was desirable
Implementation barrier	For serialised assets the identifier infrastructure exists and the barrier is institutional. For receivables the identifier infrastructure does not exist, so the identity problem must be solved before the coordination problem can be attempted

[INFERENCE, BY ANALOGY TO THE STATUTORY HISTORY]

Read against Chapter 41: three of the four candidate mechanisms in §73.7 escape the standardization-failure mechanism because the cost-bearer is the benefit-taker and no other party's agreement is required. The fourth does not, and it is the only one that answers the uniqueness question. That is the trap, stated exactly. [INFERENCE]

U. Independent challenge

Three adversarial theses, each stated at its strongest and then answered.

Thesis A: the collateral-integrity gap is invented; Article 9 and ordinary diligence already solve it. *At strength:* the filing system gives constructive notice of prior claims; first-to-file resolves contests; security agreements, field examinations, borrowing-base certificates and periodic audits do the rest; supervisory handbooks prescribe the controls; the failures are explained by lenders abandoning practices they used to follow. *Answer:* **defeated.** The failures occurred while the regime was fully operative and functioning as designed. The filing system gives notice of *lien existence*; it does not deduplicate assets across facilities, because blanket descriptions are expressly permitted and after-acquired collateral cannot be enumerated. The gap is not invented. [INFERENCE]

Thesis B: the gap is real but it is an underwriting and audit problem, not an infrastructure problem. *At strength:* red flags were present and ignored in every case; the post-failure reforms being adopted are improvements to existing practice — expanded audit and confirmation rights, direct obligor confirmation, unscheduled receivable-level audits — not new infrastructure. *Answer:* **partly proved, and the paper concedes the part that is proved.** Better diligence would have mitigated or prevented every case examined. But diligence is a single-institution capability and uniqueness is a cross-institution property; no lender's audit of its own book can see another lender's facility. [INFERENCE]

Thesis C: the gap exists but digital infrastructure is already closing it. *At strength:* a tokenised-collateral platform is under construction at a central clearing utility; the digital-asset amendments to the Code are enacted in a majority of states; vendors market cross-facility duplicate-pledge detection; proposals exist for modernised lien systems with machine-readable filings and privacy-preserving match networks. *Answer: partly proved.* The direction is real and well capitalised. But the clearing-utility platform addresses centrally cleared securities rather than warehouse lending, private credit or receivables; Article 12 addresses perfection and priority for digital assets rather than deduplication of ordinary collateral; the vendor products require originator onboarding and provide visibility only within their own networks; and the modernised-lien proposals remain proposals. As at the date of this paper, none of it covers the segments where the losses occurred. [INFERENCE]

V. The final proposition

Article 9 is primarily a legal rights, attachment, notice, perfection, priority and enforcement system. It is not an asset-integrity system. The legal system can allocate rights successfully without establishing every underlying economic fact to which those rights relate.

Classification: supported, with material narrowing.

The narrowing has two parts, and both are load-bearing:

One. Deficient single-institution diligence materially amplified the harm in every recent failure. Authoritative verification mechanisms existed in each case and were not used. The gap is therefore **primarily operational**, not primarily structural, and any account that omits this is overstating. [INFERENCE]

Two. Infrastructure aimed at parts of the gap is under active construction, and portions of it may be addressed within a horizon this paper cannot date. The proposition is strongest for receivables, warehouse-lent assets

and asset classes with no per-asset registry, and weakest for centrally cleared securities and titled goods, where authoritative records already exist. *[INFERENCE / UNKNOWN AS TO TIMING]*

What survives both narrowings, and what Chapter 73 carries into the argument: **in the asset classes that have no per-asset record — trade receivables first among them — no party is authoritative on whether a given asset has been pledged once, no single institution can establish it alone, and the mechanism that would establish it is subject to the same coordination economics that have defeated every other industry-wide initiative in this paper.**

W. Questions that require counsel

These cannot be resolved from public research and are listed in Appendix E under the legal-review heading.

1. Does a lender that did not search before advancing have a claim in tort against a debtor that double-pledged, or is the failure to search contributory?
2. Can a uniqueness-verification obligation be imposed by supervisory action — banking-agency guidance, or a securities rule for asset-backed issuance — without amending the Code?
3. Under § 9-203(b)(2), what quantum of "rights in the collateral" suffices for a pool of commingled after-acquired receivables, and where is the boundary between a right and an expectancy?
4. Does control under Article 12 create a uniqueness guarantee, or can the same controllable electronic record be subject to multiple control arrangements?
5. In a reorganisation following discovered double-pledging, how is loss allocated among innocent secured creditors each holding a facially valid filing?
6. Does a lender's reliance on a borrower's compliance certificate without independent verification constitute good faith, or can it be challenged as wilful blindness in a later avoidance action?

7. Could a lender that declined to join a voluntary uniqueness facility face liability where its non-participation was the channel through which duplication occurred?
8. Does the recent enforcement activity create a supervisory expectation of cross-facility collateral verification, such that its absence could be cited as an unsafe or unsound practice?

X. Sources and method for this appendix

Statutory conclusions rest on the text of the Uniform Commercial Code and its Official Comments, taken first and in preference to any secondary source: §§ 1-204, 9-102, 9-108, 9-109, 9-202, 9-203, 9-204, 9-210, 9-308, 9-310, 9-313, 9-318, 9-320, 9-322, 9-324, 9-406, 9-502, 9-504, 9-509, 9-516, 9-521, and Article 12 §§ 12-102 and 12-105 with the conforming amendment at § 9-102(a)(27A). One case is cited for the proposition that filing does not create a security interest: *Raleigh Industries of America, Inc. v. Tassone*, 74 Cal. App. 3d 692 (1977).

Case facts rest, in order of preference, on securities-regulator releases and litigation materials, criminal charging documents, and bankruptcy-court proceedings; where those were unavailable, on trade and legal-press reporting, which is labelled REPORTED rather than FOUND. The infrastructure survey rests on operator and regulator publications for the public systems and on vendor materials for the private ones, labelled accordingly.

Three limits apply and are stated rather than papered over. Enactment of Article 9 and of the 2022 amendments varies by state, and this appendix describes the uniform text. The securities-regulator complaint in the auto case is a pleading, not an adjudication, and every figure drawn from it is labelled ALLEGED. And the eleven-fact matrix in §E and the gap classification in §§ are analytical constructs of this paper, not terms of art drawn from the statute or the case law.

Note on sources and method

Hierarchy

Sources were weighted in the following order:

1. Primary legal texts, regulators, central banks, courts and securities filings — the Uniform Commercial Code and its Official Comments, the Federal Reserve, IMF, BIS, FSB, Bank of England, OCC, FDIC, SEC, EDGAR filings, the Second Circuit and the Delaware Court of Chancery. For the collateral analysis in Appendix F, statutory text and Official Comments were taken first and in preference to any practitioner or secondary source.
2. Independent data providers and trade bodies — S&P Global and PitchBook, Prequin, KBRA, AFME, the LSTA, LMA and APLMA, Coalition Greenwich, the Atlantic Council.
3. Peer-reviewed academic work — *Journal of Finance*, *Review of Quantitative Finance and Accounting*, software-engineering conference proceedings.
4. Tier-one financial press.
5. Law-firm client alerts and practitioner guides.
6. Vendor marketing and market-research resellers — used only for product scope, always labelled, never as evidence that a claim is true.

On what this method cannot do

Several of the paper's most important questions cannot be answered from public sources, and the paper's practice has been to say so rather than to substitute plausible reasoning. Appendix E lists them.

Three structural limits deserve statement.

Member-gated data. The trade association's settlement statistics — the market's authoritative operational measurements — are behind a membership wall. The paper's settlement figures are therefore several years old and sourced through law-firm analysis rather than from the primary study.

Vendor-dominated evidence. For the incumbent platforms, and for most of the operational-burden claims, the only available sources are the firms selling into the problem. The paper labels these and, where possible, notes when a supervisor or an independent trade body describes the same gap — which is the strongest available corroboration.

Absence of a counterfactual. No bank has publicly replaced its loan servicing core. Every argument in this paper about displacement is therefore reasoning about an event with no observed instances, and the paper's confidence is bounded accordingly.

Pleadings are not findings. Several of the collateral figures in Chapter 73 and Appendix F come from a civil complaint or an indictment. These are allegations and charges, and they are labelled as such throughout. Only the Petters matter, which ended in conviction and sentence, is in the strongest evidentiary class.

Corrections carried

The following errors are common in writing about this market and are corrected here: Loan IQ's origin (IQ Financial Systems, not Midas); ACBS's route to FIS (ALLTEL/Fidelity 2003, not SunGard); WSO's lineage (JPMorgan's FCS Corporation via Markit, not Misys); ClearPar's lineage (acquired by Markit from FIS in 2009, not built by Markit); the 2024 Finastra incident (data exfiltration, not ransomware); the two unrelated firms sharing the "AFS" name; "Advanced" rather than "Automated" Commercial Banking System; and the 2024 leveraged-loan league table ordering (Bank of America first, JPMorgan second — the reverse of the common assumption, though the ordering flips by table).

Corrections made to this paper's own earlier draft, recorded here because a paper that maintains an excluded-figures appendix owes the same discipline to itself: the aggregation of First Brands' liabilities with Tricolor's collapse size into a single loss figure; the description of that figure as "twelve-figure"; the presentation of global and US CLO estimates as if

comparable; a compressed phrasing that implied the OCC's \$400m penalty was imposed for the Revlon error; a subtitle reading "four decades" against a body that says three; the claim that a lender disagreeing with the register is "wrong regardless of the facts," which omits the manifest-error qualification; the claim that no party is authoritative about double-pledging, which is too broad given public filing regimes; and the grouping of collateral integrity with private-credit valuation as a single opportunity. To these are added the corrections made in the collateral pass: the implication that a security interest could be "perfectly perfected" over collateral that never existed, which is legally wrong because attachment fails first; the treatment of collateral existence and priority as one undifferentiated row of the authority table, when the two have different authorities and only one of them is unheld; and the absence of any acknowledgement that deficient single-institution diligence amplified the harm in every case examined. Each is listed in Appendix D or corrected in place. [See CHANGE_LOG]

BACK MATTER

ABOUT CONEKTIQ RESEARCH

Conektiq Research examines how institutions work, why they work the way they do, and what changes when artificial intelligence becomes capable of understanding and acting within complex institutional environments.

Its work combines domain expertise, primary-source investigation, machine-assisted analysis, independent challenge, and explicit evidence classification.

The objective is not to predict what AI will do. It is to determine where AI actually changes the constraints under which institutions operate.

The Architecture of Institutional Credit began as a practical test of that approach in a market the author knew well enough to challenge the system's conclusions.

ABOUT THE AUTHOR

Steve Shyn is the founder of Conektiq, where he is developing systems for institutional knowledge, governed AI work, and organizational decision-making.

Earlier in his career, he worked in institutional credit and syndicated lending.

This paper began as an experiment: could a modern AI system reconstruct a complex institutional market, distinguish evidence from inference, challenge its own conclusions, and produce work that could withstand scrutiny from people who actually know the market?

The experiment became this paper.

This paper was produced using Conektiq's institutional AI research and governance approach.

The system was used to research the market, structure evidence, distinguish fact from inference, identify contradictions, challenge conclusions, and preserve the institutional context from which the final analysis was developed.

The technology did not supply the point of view. It made it possible to investigate, challenge, and refine that point of view at a scale that would previously have required a substantial research team.

C O N E K T I Q

CONEKTIQ.COM