cXc

A TRADING NAME OF CONXCHANGE PVT LTD

RESEARCH PAPER · DISTRIBUTION SERIES · JULY 2026

From spreadsheets to a unified distribution platform

The operational, financial, and regulatory case for consolidating insurance distribution onto a single system of record.

v1.0 · Published 10 July 2026 · Classification: Public · Author: cXc Research

A companion to the cXc Distribution Series — read alongside Client money in a digital-first world, Practical applications of ML in insurance distribution, and AI in the delegated authority stack. Where those papers examine specific regulatory regimes and AI applications, this paper examines the operational substrate underneath all of them: the spreadsheet, and what happens when a distribution business outgrows it.

For CRO, COO, and Distribution leadership — at a glance

This page is designed to be lifted directly into a board or ExCo pack. It distills the paper's argument into the decisions that matter, before the evidence behind them.

The four platform decisions

- **1. Unify commission and bordereaux workflows —** stop the 5-8% manual commission error rate and the 3-5-day-per-carrier bordereau cycle from compounding every month (see Sections 4 and 5).
- **2. De-spreadsheet counterparty, client-money, and Consumer Duty evidence —** these are the layers this paper's audit and regulatory evidence identifies as highest-risk, and the ones most commonly still spreadsheet-tracked even where a core AMS is in place (see Sections 6, 7, and 8).
- **3. Sequence platform consolidation before further AI investment —** AI cannot compensate for data quality problems a spreadsheet-fragmented operating model was never designed to solve, and procuring AI first typically means paying the same integration cost repeatedly rather than once (see Section 11).
- **4. Fix the data before the system, pilot before cutover, and design governance alongside the technology —** of the two migration disciplines that separate successful consolidations from stalled ones, in that order (see Section 17).

The capability map in one line

Distribution, finance, counterparty, and evidence — one system of record, with governance running through all four. Section 14 sets this out in full.

How to use this page

Copy it directly into a board or platform-investment pack as the cover slide. The rest of the paper is the evidence base behind each of these four decisions, cross-referenced by section throughout.

Executive summary

Every paper in this series has, at some point, encountered the same artefact sitting behind the regulatory finding, the audit exception, or the AI deployment gap: a spreadsheet. Client money reconciliation breaks down into a shadow-Excel patchwork behind the official platform. Submission intake runs through a triage spreadsheet before it ever reaches a rating engine. Commission statements are matched to carrier remittances by hand, row by row, every month. This paper steps back from the specific regulatory and technical treatments elsewhere in the series to ask the more basic question: why is the spreadsheet still the default operating system for insurance distribution, and what does the evidence say about the cost of that default?

Four observations frame this paper.

One · Spreadsheet dependency is the default, not the exception

Commission reconciliation, bordereaux compilation, TOBA and delegated authority tracking, and renewal management are still run substantially in Excel across the independent broker and MGA market, at every scale from single-office agencies to platform-scale MGA groups.[1] This is not a criticism of any individual firm's discipline — it reflects a genuine gap between the pace at which distribution businesses have scaled and the pace at which their back-office tooling has kept up.

Two · The error rate is not theoretical — it is decades-documented

Raymond Panko's spreadsheet error research, built from field audits spanning nearly three decades, finds that 94% of operational spreadsheets contain at least one error, with a cell error rate around 5.2%.[8] Insurance distribution's specific complexity — multi-carrier commission structures, multi-currency bordereaux, chargebacks, and profit-share calculations — makes this a harder case than the general-business baseline the Panko research describes, not an easier one.

Three · Regulatory and technology tailwinds now converge on this exact gap

CASS 15's data-richness argument, the Consumer Duty's shift from static documentation to continuous evidence, and AI's absolute dependency on structured data all point at the

same underlying requirement: a system of record, not a spreadsheet, behind every one of these workflows. Section 13 sets out what happens, tailwind by tailwind, if a firm takes no action; each of these pressures is covered in depth elsewhere in this series, and this paper draws them together into a single operational argument.

Four · The platform question is not whether, but which capability map and which path

No firm that has read the audit, error-rate, and regulatory evidence in this paper can credibly conclude that spreadsheet-based distribution operations are a stable long-term position. The open questions are which capabilities a unified platform actually needs, how the incumbent vendor landscape maps to them, and how a firm gets from its current state to that destination without a failed migration along the way. This paper addresses all three.

The capability map, previewed

Part IV sets this out in full, with a diagram and vendor-by-vendor detail. In brief, a unified platform needs four layers, with a governance thread running through all of them:

Layer	What it does
Distribution	Submission intake, multi-carrier rating and appetite, binding and policy issuance.
Finance	Commission engine, premium accounting, bordereaux automation.
Counterparty	Counterparty master, TOBA/DUA repository, producer and agency data.
Evidence	CASS/CMAR data, continuous Consumer Duty MI, board and KRI dashboards.

How to read this paper

Part I sets out how pervasive spreadsheet dependency actually is and what the research literature says about its error rate. Part II walks through the specific operational failure modes in insurance distribution. Part III explains why this is urgent now, not a permanent background risk, and closes with what happens if a firm does

nothing. Part IV sets out the capability map in full, the vendor landscape, and the migration path. Part V closes with segment-specific implications and the cXc view.

Contents

Executive summary 4

Part I — The spreadsheet economy in insurance distribution

1. How pervasive spreadsheet dependency actually is 7
2. What the research says about spreadsheet error 9
3. Why insurance distribution is a harder case than the baseline 11

Part II — Where it breaks: the operational failure modes

4. Commission and remittance reconciliation 14
5. Bordereaux and premium flow 16
6. TOBA, DUA, and counterparty sprawl 18
7. Client money — the CASS-adjacent risk 20
8. Consumer Duty MI and the static documentation problem 21

Part III — Why this is urgent now

9. Scale outpacing tooling 24
10. CASS 15 and the PSP data-richness argument 26
11. AI's data-quality precondition 28
12. Lloyd's Blueprint Two and market-level standardisation 30
13. What happens if you do nothing 32

Part IV — What a unified platform looks like

14. The capability map 35
15. The incumbent landscape 37
16. Three migration scenarios 39
17. The migration path 42

Part V — Implications and the cXc view

18. What this means by segment 45
19. A short note on the cXc view 47
20. Glossary of working terms 48

References 49

PART I

The spreadsheet economy in insurance distribution

How pervasive it is, what the research says, and why insurance distribution is a harder case.

1. How pervasive spreadsheet dependency actually is

Before assessing the cost of spreadsheet dependency, it is worth being precise about where it actually sits in a typical distribution business. The honest answer is: almost everywhere that is not the core policy administration record itself.

Where spreadsheets do the work

- **Bordereaux compilation —** carriers require monthly, sometimes weekly, bordereaux in formats that vary by relationship — some Excel templates, others XML or CSV, others carrier-portal upload. Compiling and validating a single carrier's bordereau manually takes a finance analyst 3 to 5 business days.
- **Commission reconciliation —** matching carrier commission statements against internal records is, for a large share of the independent agency and MGA market, a manual spreadsheet exercise performed monthly, often the night before payroll.
- **Counterparty and agreement tracking —** TOBA terms, DUA authority limits, and remuneration structures are frequently tracked in a master spreadsheet rather than a structured system, particularly for firms with more than a handful of carrier and coverholder relationships.
- **Renewal and pipeline management —** even where a core AMS exists, renewal pipelines, cross-sell tracking, and producer performance reporting frequently live in a parallel spreadsheet the AMS was never configured to replace.
- **Management information —** Consumer Duty board reporting and CASS management information are, in a material share of firms, compiled monthly from multiple source systems into a spreadsheet before being presented as board evidence.

Why this persists

The persistence of spreadsheet-based operations is not a failure of ambition. MGAs operate with the underwriting and financial complexity of a small insurer but the finance headcount of a small broker — a structural mismatch that spreadsheets fill because they

are immediately available, require no procurement cycle, and can be built by the person who understands the workflow, without waiting for a system to be configured to match it.[1] That flexibility is precisely what makes spreadsheets attractive at the point a business first needs a tool, and precisely what makes them dangerous once the business has scaled past the point a single person can hold the whole model in their head.

The scale at which this becomes consequential

An MGA managing seven carrier relationships is managing seven separate premium trust flows, seven remittance schedules, and seven bordereaux formats — each historically reconciled by hand.[1] A retail broker with 80 monthly commercial submissions is generating 60 to 120 hours of underwriter data-entry time before any underwriting judgement occurs, according to McKinsey's 2024 Insurance Operations Efficiency Report, which found that data rekeying and document review consume 35 to 45% of underwriter time at firms without automated intake.[5] These are not edge cases. They describe the median operating pattern across a large share of the independent distribution market.

The test that reveals the real position

Ask a firm's finance lead how many spreadsheets are involved in producing this month's commission statement, this quarter's board MI pack, and the last CASS resolution-pack rehearsal. If the honest answer is "I'd need to check with three different people," the firm's real operating model is spreadsheet-based, regardless of what the core policy administration system's marketing material says it does.

2. What the research says about spreadsheet error

The risk spreadsheets carry is not a matter of opinion. It is one of the most thoroughly documented phenomena in applied computing research, with a research literature spanning nearly three decades and multiple independent research groups.

The Panko body of work

Raymond Panko of the University of Hawaii is the field's most-cited researcher. Reviewing thirteen field-audit studies conducted between 1995 and 2004, Panko found a weighted average cell error rate (CER) of 5.2% across 43 spreadsheets with detailed data, and a spreadsheet error rate (SER) — the percentage of spreadsheets containing at least one error — of 94% across 88 spreadsheets studied.[8] A separate, more recent field audit of 54 real-world spreadsheets, using improved detection methodology, found errors in 91% of the spreadsheets reviewed.[9]

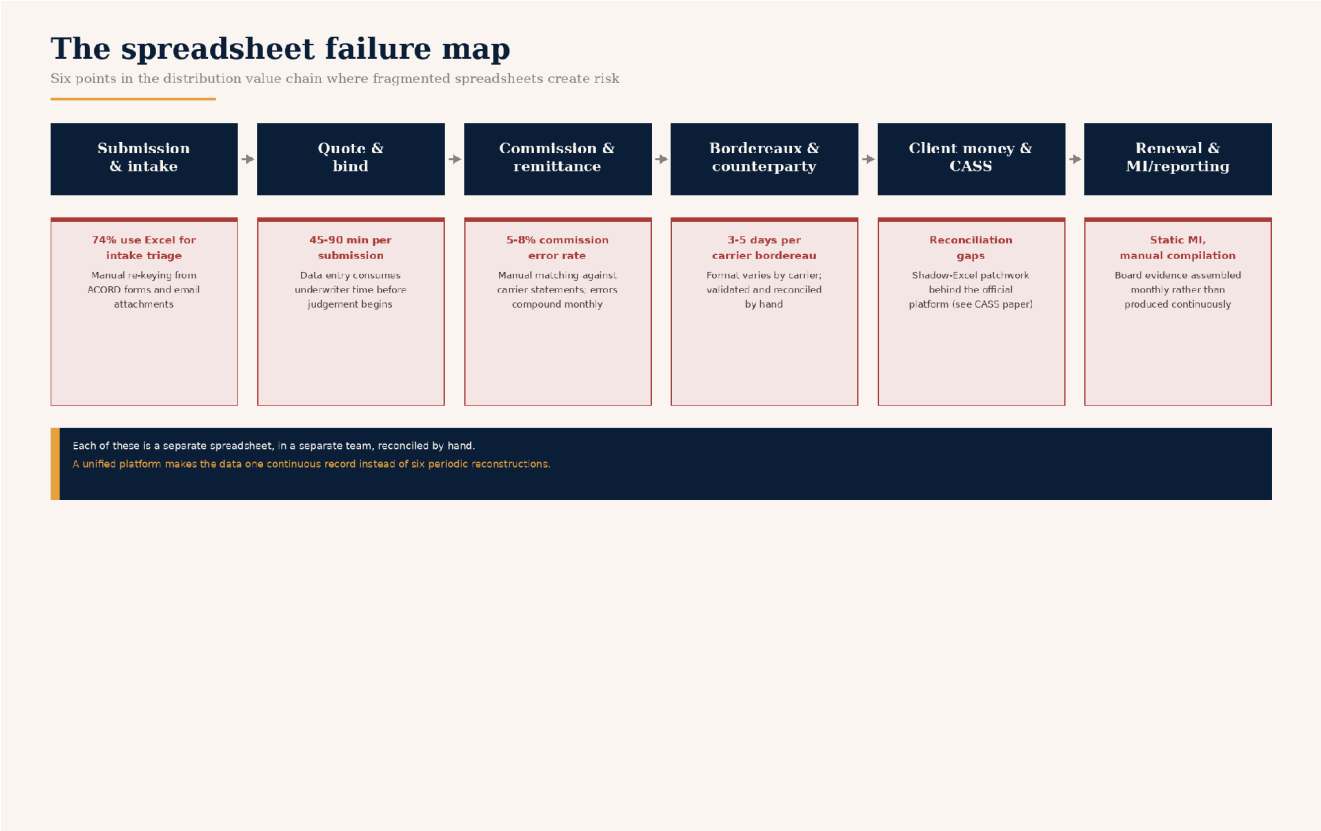

Figure 1 — Six points in the distribution value chain where spreadsheet dependency creates risk

Why this is not a matter of carelessness

Human error research across multiple domains — writing, calculation, programming — consistently finds base error rates of 1% to 5% for tasks of comparable cognitive complexity, and this pattern holds for spreadsheet development specifically.[10] This means spreadsheet errors are not a symptom of poor discipline. They are the statistically expected outcome of asking humans to perform a large volume of detailed, repetitive cognitive work without independent verification — the same finding that underlies decades of research on error rates in manual code review and manual data entry generally.

The overconfidence problem

Panko's controlled experiments found that individual spreadsheet developers estimate their own probability of error at 10% to 18%, while the actual rate of incorrect spreadsheets in the same experiments was 86%.[11] As Panko puts it, developers are "impressed with our ability to ferret out errors" while having "no idea of how many remain." This overconfidence is precisely why error-checking built into the spreadsheet itself — an additional formula, a manual cross-check — is a weak mitigation: the same cognitive blind spot that produced the original error also limits the developer's ability to catch it.

What happens when it goes wrong at scale

The JPMorgan Chase "London Whale" incident of 2012 is the canonical illustration of spreadsheet risk compounding into material financial loss. A copy-and-paste error in a risk-model spreadsheet contributed to miscalculated risk exposure, and the resulting position was a contributing factor in $6.2 billion of trading losses. The post-incident review found the error was "difficult to detect" precisely because it lived inside a working spreadsheet that appeared, on inspection, to be functioning correctly.[11]

The number that matters most

94% of spreadsheets audited in the Panko research contained at least one error. That figure has held up, with only modest variation, across nearly three decades of independent research by multiple research groups. It is not a number that improves with better training or more careful staff — the research is explicit that the error rate is a function of task complexity and volume, not individual competence.

3. Why insurance distribution is a harder case than the baseline

The Panko research describes spreadsheet error rates across general business use — budgets, forecasts, project trackers. Insurance distribution's specific operating characteristics make it a harder case than that baseline, for four structural reasons.

Reason one: multi-party commission structures

A single bound policy can generate commission splits across a producing agent, an agency principal, a wholesale broker, and (for MGA-bound business) a profit-commission calculation contingent on loss ratio performance measured over a multi-year period. Each of these splits is a formula; each formula is a candidate for the error rate the Panko research documents; and because commission calculations compound monthly, an undetected error propagates across every subsequent period until discovered.

Reason two: multi-carrier, multi-format bordereaux

Where a general-business spreadsheet has one internal format to get right, an MGA managing multiple carrier relationships is maintaining, in effect, a different spreadsheet dialect per carrier — different column structures, different aggregation rules, different currency and rounding conventions. Each format variant is a fresh opportunity for the mechanical and logic errors the spreadsheet-error literature documents, and reconciling across formats manually is exactly the kind of repetitive, detail-heavy cognitive task the human-error research shows produces the highest error rates.[9]

Reason three: regulatory retention and reconstruction requirements

A general-business budget spreadsheet with an error is corrected at the next review cycle. A commission or client-money spreadsheet with an undetected error may need to be reconstructed years later — for a CASS resolution pack, a Consumer Duty fair-value review, or an FCA supervisory request — at which point the original spreadsheet's author may have left the firm, and the workbook itself may be one of several versions with no clear record of which is authoritative.

Reason four: the audit population is not a control group

The insurance intermediary audit record gives an independent, real-world signal on exactly this question. Fewer than 1% of UK insurance intermediary CASS audits produce a clean opinion at first review, per PKF Littlejohn's published client money audit data — and the recurring findings behind that statistic (TOBA ambiguity, credit write-back errors, reconciliation cut-off failures) are precisely the class of error the spreadsheet literature predicts.[23] This is covered in full in the companion paper Client money in a digital-first world; the point for this paper is narrower: the audit evidence from actual insurance distribution firms is consistent with, not contrary to, the general spreadsheet-error research.

General business spreadsheet risk

- One internal format to maintain consistently
- Errors surface and get corrected at the next review cycle
- Reconstruction, if needed, draws on recent institutional memory
- Consequence is typically an internal management decision

Insurance distribution spreadsheet risk

- A different format per carrier, multiplied across every relationship
- Errors compound monthly through ongoing commission and premium cycles
- Reconstruction may be required years later, for a supervisory request
- Consequence can be a CASS finding, an enforcement case, or an E&O claim

The honest baseline

If 94% of general-business spreadsheets contain an error, there is no evidence-based reason to expect insurance distribution's spreadsheets to do better — and every structural reason in this section to expect them to do worse. The multi-party, multi-carrier, multi-year nature of insurance distribution's core financial workflows is precisely the profile the human-error research associates with the highest error rates.

PART II

Where it breaks: the operational failure modes

Commission, bordereaux, counterparty sprawl, client money, and Consumer Duty MI.

4. Commission and remittance reconciliation

Commission reconciliation — checking that what a carrier actually pays matches what the firm's own records say it should have earned — is the single most commonly cited spreadsheet failure point in the operational literature, and the one with the most direct revenue consequence.

5-8%

the manual commission error rate reported across MGA and agency finance operations, falling to under 1% with automated matching.

Peakflo, insurance MGA finance automation research, 2026

Why manual reconciliation produces this error rate

The mechanics of manual commission reconciliation are exactly the repetitive, detail-heavy task the spreadsheet-error literature associates with the highest error rates: matching hundreds or thousands of statement rows against internal ledger entries, across producer splits, override structures, and chargeback timing, with no independent verification beyond a second look by the same or another manual reviewer.[38] A single mistyped formula or a broken reference in the matching spreadsheet does not fail loudly — it silently misallocates commission until someone notices a producer complaint or a carrier underpayment large enough to be visible against noise.

Where the money actually goes missing

- **Direct bill policies —** carrier statements differ from internal records due to mid-term endorsements, cancellations, and timing differences between when a change is recorded internally and when the carrier's system reflects it.
- **Collection-model timing —** commission timing differs by payment model (annual, semi-annual, monthly), and spreadsheets tracking multiple models for different carriers are a common source of premature or missed accrual.

- **Contingent and profit commission —** where a commission structure is contingent on portfolio loss ratio performance measured over a multi-year period, the calculation is complex enough that a single spreadsheet error can misstate a material sum without any single row looking obviously wrong.
- **Unrecovered underpayments —** underpayments that go undetected because no one is specifically looking for them are, by definition, invisible until a full audit — and a full audit is exactly the exercise most firms are trying to avoid by using a spreadsheet in the first place.

The compounding effect

Unlike a one-off spreadsheet error in a budget forecast, a commission reconciliation error recurs every period until it is found. An error in a producer's split percentage, introduced once and never corrected, does not cost the firm once — it costs the firm every month the affected policies remain on the books, and it costs the producer the same amount in underpayment, with the corresponding risk to producer trust and retention that follows from discovering the error late.

What automated reconciliation actually changes

The reported effect of moving commission reconciliation from spreadsheet-based matching to automated matching against carrier statement data is not just an error-rate improvement — it is a change in when errors are caught. A firm that catches a misallocation within the same commission cycle can correct it before it compounds. A firm that catches it at the next full audit is correcting months or years of accumulated misallocation at once, with all of the producer-trust and cash-flow consequences that implies.

5. Bordereaux and premium flow

Bordereaux — the periodic reports summarising premium, exposure, and claims data between an MGA or coverholder and its carrier partners — sit at the centre of the delegated authority relationship, and their manual compilation is one of the most time-intensive spreadsheet-dependent processes in the distribution value chain.

3-5 days

the time a finance analyst typically needs to compile, validate, and reconcile a single carrier's monthly bordereau by hand.

IOFM, 2024, cited in Peakflo MGA finance automation research

Why bordereaux compilation is uniquely spreadsheet-dependent

Every premium dollar an MGA collects flows through a specific path: collected from the insured or retail agent, held in trust for the carrier rather than as an MGA operating asset, less the MGA's commission, remitted to the carrier on a specified schedule. An MGA with seven carrier relationships manages seven separate premium flows with different trust accounting requirements, different remittance schedules, and different carrier reporting formats.[1] Each carrier's bordereau format varies — some require Excel templates, others XML or CSV feeds, others direct submission through carrier portals — which means an MGA's bordereaux compilation process is, in practice, a different manual translation exercise for every carrier relationship.

What this does to the monthly close cycle

Finance teams that automate bordereaux, premium reconciliation, and commission workflows report a 50 to 65% reduction in month-end close-cycle time. For an MGA whose close currently runs 14 business days, that recovery of 7 to 9 business days per month is not simply an efficiency gain — it is capacity that would otherwise be spent on manual

reconciliation being redirected to financial analysis, carrier relationship management, and growth planning.[1]

The regulatory dimension

Bordereaux inaccuracy is not merely an internal inefficiency. Where a carrier's Department of Insurance examination, or a UK CASS audit, tests the underlying premium trust accounting behind the bordereaux a firm has submitted, the cost of a single adverse examination finding — legal fees, remediation work, potential penalties — typically exceeds the full annual cost of a finance automation platform that would have prevented it.[1] This mirrors the finding in this series' Client money paper: audit qualification risk and remediation cost are the largest, and most avoidable, cost associated with manual reconciliation.[23]

The Lloyd's-specific dimension

Delegated authority now represents just under 40% of business at Lloyd's, with a global network of more than 4,000 approved coverholders — each producing bordereaux into the market's central settlement infrastructure. Section 12 of this paper covers the Lloyd's Blueprint Two programme's implications for bordereaux standardisation; the CASS paper's Section 12 covers the Lloyd's-specific CASS considerations in full.

6. TOBA, DUA, and counterparty sprawl

Terms of Business Agreements and Delegated Underwriting Authority Agreements are the contractual foundation of every intermediary relationship a distribution business holds — and they are frequently tracked in exactly the tool least suited to contractual and authority-boundary precision: a spreadsheet master list.

What a counterparty spreadsheet typically tries to hold

- **Per-counterparty TOBA terms —** trust status, risk-transfer terms, remuneration structure, and any bespoke clauses negotiated outside the standard template.
- **Per-binder DUA authority limits —** premium bands, class-of-business restrictions, geographic scope, and referral thresholds, for every binder a firm operates under.
- **Version and amendment history —** which version of a TOBA or DUA is currently in force, when it was last reviewed, and what changed between versions.
- **Signatory and countersignature tracking —** who signed, when, and whether the counterparty's own systems reflect the same terms the firm's spreadsheet records.

Where this breaks down in practice

A spreadsheet-based counterparty register cannot enforce anything. It can only record what a human remembers to update, at the point a human remembers to update it. Authority creep — where a bind is accepted that falls outside the DUA's actual terms — is not caught by the spreadsheet at the moment of the bind; it is caught, if at all, at the next periodic audit, by which point the exposure has already been written.[9]

This is precisely the failure mode the delegated authority governance literature has identified: authority logic has to become machine-readable once a business scales past the point a single underwriter can hold every binder's terms in working memory, because a spreadsheet-and-referral-email model executes what an individual remembers, not what the binder actually allows.[9] The MGA companion paper in this series, AI in the delegated authority stack, covers the authority-as-code architecture that addresses this directly.

Multi-carrier commission setup as a specific instance

Standard policy administration systems were built with a simple, linear relationship in mind: one company writing insurance on one core system. Modern MGAs managing multiple carrier programmes simultaneously find that inflexible commission setup — a rating and administration system that cannot adapt to varying carrier agreements — pushes commission logic back out into a spreadsheet layered on top of the system of record, precisely the shadow-IT pattern this paper is describing.[42]

The three-question counterparty test

For any specific counterparty relationship, can the firm answer three questions within minutes rather than hours? First: what are this counterparty's current TOBA or DUA terms, and when were they last reviewed? Second: has every bind under this relationship in the last quarter been within those terms? Third: if this relationship ended tomorrow, could the firm produce a complete, current record of every open exposure under it? A spreadsheet-based counterparty register fails this test not because the data does not exist somewhere, but because assembling the answer requires a person, not a query.

7. Client money — the CASS-adjacent risk

Client money handling is covered in full in this series' companion paper, Client money in a digital-first world — a 99-page treatment of CASS 5, CASS 7, the new CASS 15 safeguarding regime, and the operational architecture that satisfies the trust obligation at scale. This section is deliberately short: its purpose is to place the spreadsheet-dependency argument of this paper alongside that paper's evidence, not to repeat it.

The single-sentence connection

Fewer than 1% of UK insurance intermediary CASS audits produce a clean opinion at first review. The recurring findings behind that statistic — TOBA ambiguity, credit write-back errors, reconciliation cut-off failures, claims-money delay — are, without exception, failure modes a shadow-Excel patchwork behind the official platform produces reliably and a properly designed reconciliation architecture is built to prevent.[23]

Where the two papers meet

- **Failure modes —** the CASS paper's risk-and-controls map (Section 18) names credit write-back, TOBA ambiguity, and reconciliation cut-off discipline as the recurring failure modes; every one of them is a spreadsheet failure mode in the general sense this paper documents in Sections 2 and 3.
- **The capability map —** the CASS paper's platform capability map (Section 25) sets out the reconciliation engine, counterparty repository, and CMAR/CASS metrics mart that a unified platform needs for client money specifically; Part IV of this paper generalises that capability map across the whole distribution operating model, not just the CASS-regulated slice of it.
- **The evidence argument —** both papers converge on the same finding: a platform that produces continuous, system-generated evidence is fundamentally different from — not just faster than — one that reconstructs evidence from spreadsheets at the point a supervisor or auditor asks for it.

Read this alongside the CASS paper

This paper's Sections 2 and 3 supply the general evidentiary base — decades of spreadsheet-error research. The CASS paper supplies the regulatory detail and the CASS-specific capability map. Read together, not as substitutes for each other.

8. Consumer Duty MI and the static documentation problem

The FCA's supervisory posture on Consumer Duty evidence has shifted in a way that makes spreadsheet-based management information a specific, named liability rather than a generic inefficiency.

From process evidence to outcome evidence

The FCA's review of Year 2 Consumer Duty board reports found that firms are, in general, good at recording what they did — process evidence — but comparatively weak on evidence of what happened to the customer as a result — outcome evidence. Reports were, in the regulator's own description, "long on narrative but short on actionable MI," with boards asked to sign off on compliance assertions they had limited ability to independently test.[81]

Why spreadsheet-compiled MI produces exactly this weakness

Management information assembled monthly from multiple source systems into a spreadsheet is, definitionally, a snapshot — a reconstruction of a point in time, built after the fact, from whatever data happened to be extractable at the point someone sat down to build the report. It cannot show a board the same-day exception status a live system produces, and it structurally discourages the kind of granular, interaction-level evidence the FCA's good-practice examples describe: firms that could link a complaints trend to a specific technical incident, or a call-abandonment spike to a specific resourcing change, were the firms whose underlying data infrastructure supported that level of granularity as a matter of course, not the firms compiling a monthly spreadsheet summary.

What good practice actually requires

The FCA's own good-practice examples are explicit about the standard: firms doing best at monitoring outcomes use good-quality data, have clear escalation processes to identify poor outcomes, and can show what they did when they found one and whether the action worked. Every one of these is a live-system capability. None of them is achievable, at scale, from a monthly spreadsheet compilation exercise.[82]

The fair-value connection

Price-and-value evidence — demonstrating a "reasonable relationship between the price charged and profitability" — is one of the specific outcomes the Consumer Duty requires firms to evidence continuously, not retrospectively construct when asked. A firm whose pricing, commission, and remuneration data all live in separate spreadsheets, reconciled against each other only at the point a fair-value review is due, is manufacturing exactly the kind of static, defensive documentation the FCA's Year 2 review flagged as inadequate.

The board test

Ask whether the firm's board, at its next meeting, could see same-day exception data on a Consumer Duty outcome of its choosing — or whether the answer is "we can pull that together for the next quarterly pack." The gap between those two answers is the gap between a platform-supported evidence base and a spreadsheet-compiled one, and it is precisely the gap the FCA's Year 2 board report review identified across the market.

PART III

Why this is urgent now

Scale, CASS 15, AI's data-quality precondition, Lloyd's Blueprint Two, and what happens if you do nothing.

9. Scale outpacing tooling

Spreadsheet dependency has always carried the risks documented in Part I and Part II. What makes 2026 a specific inflection point, rather than a permanent background condition, is that the scale of the distribution market has grown faster than the tooling most firms run on.

$1.07t rn

US P&C direct written premium — competing for that scale on submission workflows built for 2010 is a compounding disadvantage.

Insurance Information Institute, 2025 Fact Book

The MGA growth curve specifically

US MGA premium grew from $47 billion in 2020 to $97 billion in 2024 — roughly a 14% compound annual growth rate, well ahead of the broader insurance market's trajectory. Delegated authority now exceeds 40% of business at Lloyd's.[1] Every one of those growing MGAs is managing more carrier relationships, more bordereaux formats, and more commission structures than it was five years ago — and a spreadsheet-based operating model that was merely inefficient at $20 million in premium becomes structurally unsafe well before $200 million.

M&A consolidation compounds the problem

MGA and MGU groups at platform scale are transacting at 15 to 22 times adjusted EBITDA, against 10 to 14 times at the lower-middle-market tier — private equity-backed consolidation is actively assembling groups of previously independent agencies and MGAs.[6] Every acquisition in that consolidation pattern brings its own spreadsheet estate, its own counterparty tracking conventions, and its own commission logic into the combined group — multiplying, rather than resolving, the fragmentation this paper

describes, unless the acquirer has a deliberate data-consolidation plan in place before the deal closes.

The talent constraint

A commercial-lines underwriter is the most expensive seat in distribution, and the supply of senior commercial underwriters is constrained. Every hour spent on manual data entry and spreadsheet reconciliation is an hour not spent on the underwriting judgement that actually requires that expensive seat — which means spreadsheet dependency is not just an error-rate risk, it is a capacity constraint on growth that compounds as the talent market tightens.

The scaling test

Ask whether the firm's operating model today would still work, without modification, at twice its current premium volume. If the honest answer involves "we'd need to hire more people to do the same manual reconciliation," the firm is scaling headcount against a fixed unit cost of manual work, rather than scaling volume against a largely fixed platform cost — a fundamentally worse economic position as growth continues.

10. CASS 15 and the PSP data-richness argument

This series' Client money paper covers CASS 15 — the FCA's new safeguarding regime for payment and e-money institutions, in force since 7 May 2026 — in full detail. This section draws out the specific implication for the spreadsheet-to-platform argument.

What CASS 15 changes on the PSP side

Every PSP and EMI in a firm's premium-collection chain is now producing daily reconciliation, individual-customer-level records, and structured resolution-pack data as a matter of regulatory obligation, not as an optional service enhancement.[23] This is a one-directional change: the data richness now arriving from PSP and EMI counterparties is not something a firm can decline to receive, and it is arriving at an event-driven cadence that a spreadsheet-based reconciliation process was never designed to consume.

The specific missed opportunity

A firm whose own reconciliation architecture is still spreadsheet-based and batch-oriented is now paying the integration and safeguarding tax on the PSP side of the relationship — which it cannot avoid, since the PSP's obligations are independent of the firm's own systems — while remaining unable to use the better data now available, because nothing in its own operating model is built to consume it at the pace it now arrives. The firm is, in effect, funding a data quality improvement it cannot benefit from.

Why this argues for acting inside the current window

CASS 15's Supplementary Regime is now in force; the Post-Repeal full statutory trust framework remains subject to consultation.[23] That means PSP and EMI data formats and integration patterns are still stabilising rather than fully settled. Firms that build the consuming architecture now, while the pattern is still forming, face an easier integration than firms that wait until the PSP relationship and data format are fully established and then retrofit — a variant of the same build-versus-retrofit argument that runs through every section of this paper.

The connective thread

CASS 15 is a payments-regulation story, not a distribution-platform story — but its practical effect is to make richer, more frequent data available to every firm using a PSP for premium collection. A spreadsheet-based reconciliation process cannot use that data at the cadence it now arrives. A platform built around continuous reconciliation can. The regulatory change did not create this gap; it made the gap between the two operating models measurably wider.

11. AI's data-quality precondition

This series' AI paper, Practical applications of ML in insurance distribution, documents where machine learning and generative AI now produce measured, repeatable returns in distribution — submission intake, lead scoring, retention modelling, algorithmic follow. Every one of those applications has an unstated precondition this paper makes explicit: none of them work well against spreadsheet-fragmented data.

Garbage in, garbage out — the specific mechanics

Submission intake AI's published operating characteristics — 60 to 90% reduction in intake time, 97 to 98% first-pass accuracy — are achieved against a well-instrumented, structured event pipeline, not against data first captured in a spreadsheet and later exported.[24] AI-augmented commission and bordereaux processing depends on the same precondition: a model that reconciles carrier statement data against internal records needs both sides of that comparison in structured form, and a spreadsheet on either side reintroduces exactly the manual translation step the AI application was meant to remove.

Where firms get this backward

A recurring pattern in AI deployment failures — covered in the parent AI paper's treatment of failure modes — is a firm procuring an AI capability to sit on top of its existing spreadsheet-based operating model, expecting the AI to compensate for data quality problems the underlying architecture was never designed to solve. This produces the worst of both outcomes: the AI license cost is paid, and the underlying data quality problem persists, frequently made harder to diagnose because the AI layer obscures where in the pipeline a given error originated.

The sequencing implication

The build-versus-buy integration-tax pattern documented in the AI paper — a representative £500k vendor licence producing £1.2m of first-year integration cost — is, in large part, the cost of building the structured-data foundation an AI capability needs where a spreadsheet-fragmented operating model does not already provide it.[24] Firms that

address the platform-consolidation question in this paper before procuring AI capability are addressing that integration cost once, as infrastructure; firms that procure AI first are often paying the same integration cost per AI capability, repeatedly.

The sequencing rule

Platform consolidation is not a prerequisite for every AI use case — document classification at the point of intake works reasonably well against semi-structured inputs. But any AI application that reconciles two data sources against each other — commission matching, bordereaux validation, cross-sell against a clean customer record — needs the structured-data foundation this paper describes first.

12. Lloyd's Blueprint Two and market-level standardisation

The clearest signal that the wider market agrees spreadsheet-based operations have reached their limit is Lloyd's own multi-year investment in replacing manual, spreadsheet-mediated market processes with standardised digital infrastructure.

What Blueprint Two is actually doing

Lloyd's Blueprint Two programme continues to invest in coverholder onboarding, bordereaux processing, and market settlement infrastructure, with the explicit goal of moving bordereaux and premium settlement away from bespoke, per-relationship spreadsheet formats toward standardised, structured data exchange across the market.[23] This is a market-level acknowledgement of the same problem this paper documents at firm level: bordereaux compiled and exchanged as spreadsheets do not scale, and the fix is standardised, structured data, not better spreadsheet templates.

What this means for firms operating in the London Market

- **Early movers benefit twice —** coverholders and brokers whose own systems already produce structured, API-consumable data will integrate with Blueprint Two's standardised infrastructure far more easily than firms whose bordereaux process starts from a spreadsheet and requires a manual conversion step before it can be submitted in the new format.
- **The standardisation cadence will keep moving —** as Blueprint Two's data standards mature, market-level pressure to conform will increase; firms with a unified platform behind their bordereaux process will adapt to each standardisation milestone as a configuration change, while firms with a spreadsheet-based process will face a re-templating exercise each time.
- **Governance maturity is accelerating alongside infrastructure —** the London Market Association's own April 2026 survey found AI adoption across Lloyd's had more than doubled in twelve months, with 93% of respondents building formal

governance frameworks — market-level digital maturity is accelerating in parallel with Blueprint Two, not independently of it.

The market is not waiting

A firm can reasonably debate the pace of its own platform consolidation. It cannot reasonably assume the wider market infrastructure it depends on will continue to accommodate spreadsheet-based bordereaux and premium settlement indefinitely. Blueprint Two is a multi-year programme already under way; the firms that treat platform consolidation as optional are making a bet against the direction Lloyd's has already committed to.

13. What happens if you do nothing

Parts I through III have each made a version of the same argument from a different direction: spreadsheet dependency is measurably risky, and four independent tailwinds now converge to make that risk urgent rather than merely chronic. Before Part IV sets out what to build, this section makes the converse case explicit — what a firm should expect if it takes no action on any of the four tailwinds already in motion.

Tailwind	If a firm does nothing
CASS 15 (Section 10)	PSP and EMI counterparties produce richer, event-driven data regardless of the firm's own readiness. A spreadsheet-based reconciliation process pays the integration and safeguarding tax on the PSP side while remaining unable to use the better data now available on its own side.
Consumer Duty MI (Section 8)	The FCA's shift from static documentation to continuous outcome evidence continues independent of any single firm's response. A firm still compiling MI monthly from spreadsheets is producing exactly the "long on narrative, short on actionable MI" evidence the FCA's Year 2 board report review already flagged as inadequate — at the next review cycle, not some future one.
AI adoption (Section 11)	AI capability procured on top of fragmented, spreadsheet-based data continues to underperform or fail outright, and the integration tax gets paid repeatedly, once per AI use case, rather than once as shared infrastructure.
Lloyd's Blueprint Two (Section 12)	Market-level bordereaux standardisation proceeds on Lloyd's own timeline regardless of any single firm's readiness. Firms whose bordereaux process starts from a spreadsheet face a re-templating exercise at every standardisation milestone; firms with structured data behind the process absorb each milestone as configuration.
Scale and M&A consolidation (Section 9)	Growth and acquisition continue to add carrier relationships, commission structures, and spreadsheet estates faster than a manual operating model can absorb them, regardless of whether the firm has made an explicit platform decision.

The point of this table

None of these five rows describes a risk that stays constant while a firm decides whether to act. Every one of them is already moving — a regulator's supervisory cadence, a market infrastructure programme, a competitor's AI deployment, an acquisition pipeline. "No action" is not a neutral holding position on any of the four tailwinds this paper has documented. It is a bet that none of them will arrive before the firm is ready, placed simultaneously against all five rows in this table.

PART IV

What a unified platform looks like

The capability map, the incumbent landscape, and the migration path.

14. The capability map

Parts I through III established why spreadsheet dependency is a structural risk and why regulatory and market pressure now converge on fixing it. This section sets out, concretely, what a unified distribution platform needs to do — a capability map any credible build, buy, or configure decision should be checked against.

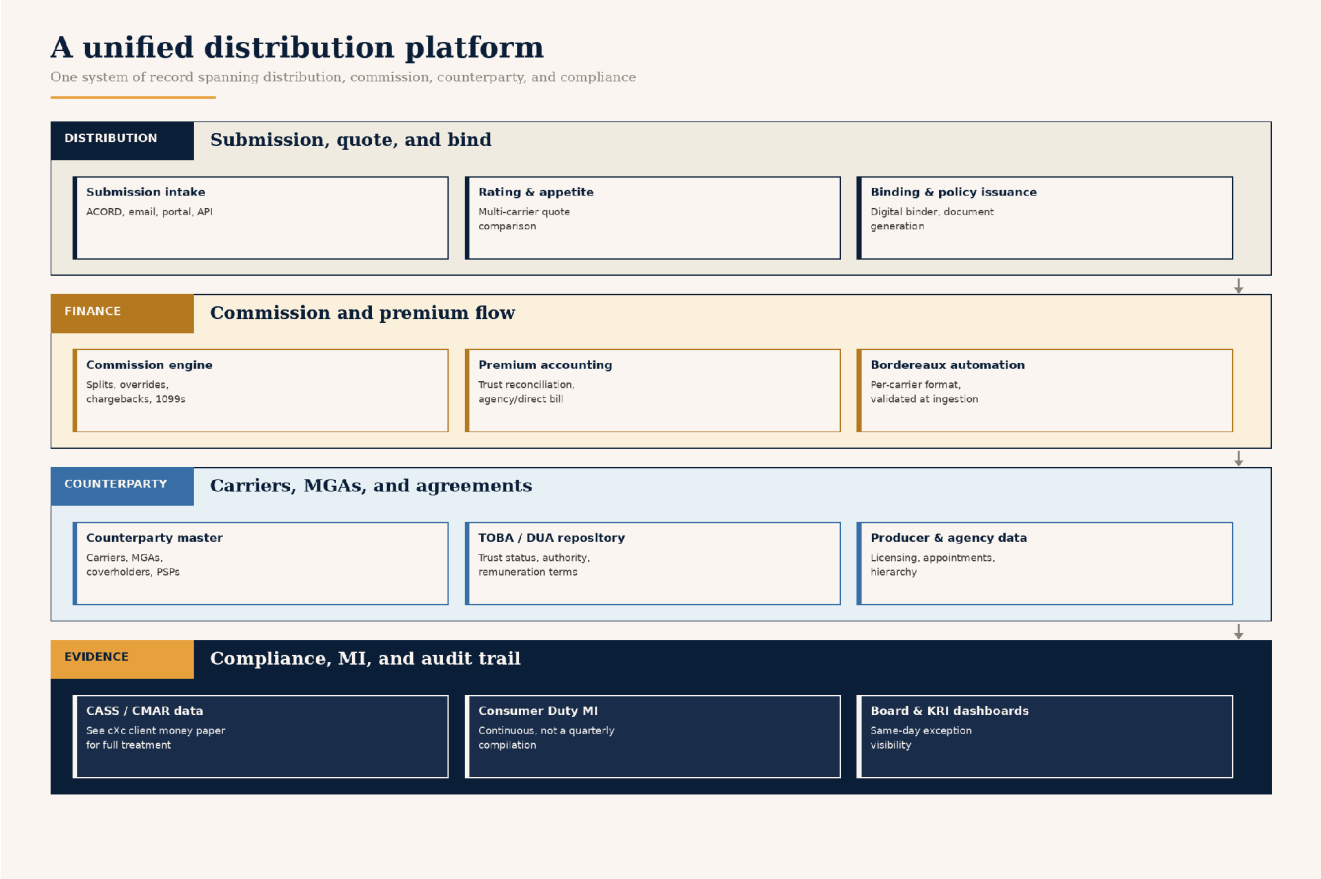

Figure 2 — A unified distribution platform: four layers, one system of record

The four layers

Distribution — submission, quote, and bind

Submission intake across channels (ACORD, email, portal, API), multi-carrier rating and appetite matching, and binding with automatic policy issuance. This is the layer with the deepest AI penetration and the most mature vendor landscape, covered in detail in the AI paper's treatment of submission intake and quote-and-bind.

Finance — commission and premium flow

A commission engine handling splits, overrides, chargebacks, and statutory reporting; premium accounting spanning trust reconciliation and agency-bill and direct-bill workflows; bordereaux automation that validates against carrier formats at ingestion rather than reconciling after the fact. This is the layer where Sections 4 and 5 of this paper locate the most acute spreadsheet-dependency risk.

Counterparty — carriers, MGAs, and agreements

A single deduplicated counterparty master spanning carriers, MGAs, coverholders, and PSPs; a TOBA and DUA repository holding trust status, authority terms, and remuneration structure as structured data rather than PDF archive; and producer and agency data covering licensing, appointments, and hierarchy. This is the layer Section 6 identifies as most commonly still spreadsheet-tracked even where the other layers have moved to a proper system.

Evidence — compliance, MI, and audit trail

CASS and CMAR data (see the Client money paper's full treatment), continuous Consumer Duty management information rather than a quarterly compilation exercise, and board and KRI dashboards with same-day exception visibility. This is the layer Section 8 identifies as the FCA's specific current supervisory focus.[23]

The governance thread running through all four layers

RACI-driven approvals, change management for pricing and commission logic, and audit response sit across all four layers rather than inside any one of them — the same governance discipline the CASS paper's Section 26 sets out for client money specifically, generalised here across the whole operating model.

Using this as a specification, not a feature checklist

Any vendor evaluation, build specification, or RFP should be checked against these four layers explicitly, not against a generic feature list. A platform strong on distribution and finance but weak on the evidence layer will win a sales demo and fail

a CASS audit. A platform strong on all four layers but with no governance thread running through them will produce good data that nobody is specifically accountable for acting on.

15. The incumbent landscape

The vendor landscape serving insurance distribution is broad, regionally fragmented, and unevenly mature across the four capability layers set out in Section 14. This section maps the major incumbents against where they actually sit.

The UK market

Acturis is the dominant platform in the UK broker and MGA market by a wide margin: over 850 broker, insurer, and MGA partners transact more than £18.5 billion of premium annually on the platform, with more than 32,000 UK users and over 100,000 users globally across the Acturis Group's UK, German, French, and Canadian operations.[19] Acturis is strongest on the distribution layer — quoting, binding, and e-trading — and has been extending into accounts and document management; the finance and counterparty layers are typically where UK Acturis users still run supplementary spreadsheet processes alongside the core platform.

Open GI holds a strong position in UK personal-lines and SME commercial broking specifically, with a long-established footprint among thousands of UK broking firms, though a narrower fit for larger and more specialised commercial operations.[21]

The US market

Applied Epic, from Applied Systems, is the dominant agency management system for larger independent agencies and regional brokers with complex commercial-lines books, with the deepest carrier real-time integration library of any competitor and IVANS integration for automatic policy download.[17] Vertafore's AMS360 and Sagitta platforms hold a comparably strong position, particularly for commission and channel-hierarchy management. Both are AMS-first platforms — strongest on the distribution layer, with commission and counterparty capability typically added through integration modules or third-party add-ons rather than native design.

The MGA-specialist tier

A newer generation of platforms — BindHQ, Insly, Turris, AgentSync among them — target the MGA and program-administrator segment specifically, several explicitly positioning against the fragmented AMS-plus-spreadsheet pattern this paper describes. Turris reports migrating customer data from legacy portals, spreadsheets, and manual processes other vendors will not touch, and reports an average decrease in wait-time-to-first-bind of over 90% for its distribution-management customers.[18] These platforms are generally strongest on the counterparty and distribution layers, with finance and evidence capability maturity varying significantly by vendor and worth specific diligence.

The enterprise-scale tier

Guidewire, Duck Creek, Sapiens, and Insurity serve carrier-side policy administration primarily, with distribution-side capability typically extended through partner integrations rather than native distribution-management functionality. These platforms fit large carriers and large MGA groups with dedicated implementation teams; they are generally not the right starting point for a mid-market broker or MGA's platform consolidation question.[21]

Platform tier	Strongest layer	Typical fit	Where spreadsheets typically persist
Acturis / Open GI (UK)	Distribution	UK brokers and MGAs of most sizes	Finance, counterparty
Applied Epic / Vertafore (US)	Distribution	US independent agencies, regional brokers	Finance, counterparty, evidence
MGA-specialist (BindHQ, Insly, Turris)	Counterparty, distribution	MGAs and program administrators	Evidence, varies by vendor
Enterprise PAS (Guidewire, Duck Creek, Sapiens)	Carrier-side policy admin	Large carriers, large MGA groups	Distribution-side integration gaps

The diligence question this table implies

No incumbent platform in this landscape covers all four capability-map layers natively and completely. The practical question for any firm is not "which platform does everything" but "which platform is strongest where our current spreadsheet dependency is worst, and what does the integration or complementary-tool plan look like for the remaining layers."

16. Three migration scenarios

The build-versus-buy question is easier to reason about as three distinct operating scenarios, each with a materially different risk and growth profile, rather than as a single technology decision.

Scenario A — Continue with the current fragmented estate

Core AMS or policy administration system for the distribution layer, spreadsheets for finance and counterparty tracking, a monthly compilation exercise for evidence and MI. This is the default state for firms that have grown without a deliberate platform-consolidation decision. The likely pattern: commission error rates remain in the 5 to 8% range documented in Section 4; bordereaux compilation continues to consume 3 to 5 days per carrier relationship; and every new carrier relationship, every new distribution channel, and every new regulatory expectation adds further manual work rather than plugging into an existing pattern.

Scenario B — Consolidate the finance and counterparty layers onto point solutions

Keep the existing AMS for distribution; add a dedicated commission-and-bordereaux automation tool and a separate counterparty or compliance tracking tool. This is a meaningful improvement over Scenario A — the point solutions in Section 15's MGA-specialist tier are specifically built for this problem — but it introduces a new integration burden: data now has to move correctly between three or four systems rather than being fragmented across spreadsheets within one. The likely pattern: material error-rate and cycle-time improvement in the layers addressed, with the evidence layer (continuous Consumer Duty MI, board dashboards) often remaining a manual compilation exercise unless specifically addressed.

Scenario C — A genuinely unified platform across all four layers

Distribution, finance, counterparty, and evidence data live in one system of record, with the governance thread from Section 14 running through all four. New carrier

relationships, new distribution channels, and new regulatory reporting requirements are absorbed as configuration rather than triggering a new integration project. The likely pattern: commission and bordereaux processing generate their own audit trail as a by-product of normal operation; CMAR, CASS, and Consumer Duty evidence are queries against live data rather than monthly reconstructions; and the firm's capacity to take on new carrier relationships or new distribution channels scales with configuration effort rather than headcount.

What Scenario A costs, compounding

- 5-8% commission error rate, recurring every cycle
- 3-5 days per carrier bordereau, every month
- Audit qualification risk stays elevated
- Every new channel adds bespoke manual work

What Scenario C changes structurally

- Commission matching automated against carrier data at source
- Bordereaux validated at ingestion, not reconciled after the fact
- CMAR and Consumer Duty evidence produced as a live-data query
- New channels onboard against a defined integration pattern

Indicative impact areas

None of the figures in this paper were collected as a firm-specific business case, and none should be treated as one. But the third-party research already cited in Parts I and II — not new figures introduced here — can be read together as the order-of-magnitude shape of what platform consolidation typically changes, for a COO or CFO scoping the size of the opportunity before commissioning a firm-specific study.

Impact area	Indicative direction, from cited research
Commission accuracy	5-8% manual error rate toward under 1% automated (Section 4)
Bordereaux cycle time	3-5 days per carrier toward hours, at ingestion (Section 5)

Impact area	Indicative direction, from cited research
Month-end close cycle	50-65% reduction reported where reconciliation is automated (Section 5)
Underwriter data-entry time	35-45% of underwriter time currently consumed at unautomated MGAs (Section 1)
First-year integration cost	A representative £500k licence produces roughly £1.2m of integration cost (Section 11, from the AI paper)

Read this table as a scoping input, not an answer

Every row is a range from cited third-party or vendor research, not a cXc estimate and not a substitute for a firm-specific business case. The value of the table is directional: it tells a CFO roughly where the size of the opportunity sits before commissioning the analysis that would produce a firm-specific number.

The honest self-assessment this table demands

Most firms reading this paper are not honestly at Scenario A or Scenario C — they are somewhere in Scenario B, with some layers consolidated and others still spreadsheet-dependent. The value of this three-scenario framing is not to force a firm into a category; it is to make explicit which of the four capability-map layers from Section 14 the firm's current operating model actually addresses, and which it does not.

17. The migration path

Firms that have read Sections 13 through 15 and concluded that Scenario C is the right destination still face a genuine practical question: how does a firm actually get there without a failed, over-scoped migration programme along the way. This section sets out the sequencing pattern that distinguishes successful consolidations from stalled ones.

Start with the layer causing the most damage, not the easiest layer

The instinctive sequencing choice is often to start with whichever layer is easiest to migrate — frequently the distribution layer, since AMS migrations are a well-understood, well-supported category. But the evidence in Parts I and II points at finance and evidence as the layers producing the most measurable cost: commission error, bordereaux compilation time, and CASS-adjacent audit risk. A migration sequenced by damage reduction, rather than by ease, addresses Sections 4, 5, 7, and 8's findings directly rather than deferring them.

Migrate data quality before migrating systems

A recurring failure pattern in platform migrations is moving fragmented, error-prone spreadsheet data directly into a new system without first correcting it — which produces a new platform running on the same bad data, at greater expense. The disciplined sequence is: standardise and clean the underlying data model first, even before any new system is selected; then migrate into the new platform against a data set that is already fit for purpose.[23]

Run a defined pilot before a full cutover

A single carrier relationship, a single product line, or a single office is a reasonable pilot scope for testing a new commission or bordereaux workflow before committing the whole book to it. The pilot's purpose is not just technical validation — it is building the internal evidence base (error rates before and after, cycle-time before and after) that makes the case for full rollout to a board or investment committee that has not yet seen the numbers this paper cites.

Treat governance as part of the migration, not a follow-on project

The governance thread running through Section 14's capability map — RACI, change management, audit response — has to be designed alongside the technology, not retrofitted after go-live. A platform migration that solves the data problem but leaves the accountability question unanswered produces a better-quality spreadsheet problem: good data that nobody is specifically responsible for acting on.[23]

Expect, and budget for, the integration tax

The AI paper's representative integration-tax pattern — a £500k vendor licence producing £1.2m of first-year integration cost — is a reasonable planning assumption for platform consolidation generally, not just for AI-specific procurement. Firms that budget for this figure from the outset make better vendor and sequencing decisions than firms that discover it partway through a migration that was scoped against the licence cost alone.[24]

The sequencing rule that matters most

Fix the data before the system, pilot before the cutover, and design the governance alongside the technology rather than after it. Every documented platform-migration failure pattern in this market traces back to skipping one of these three disciplines — usually in the pursuit of a faster timeline that, in practice, produces a slower and more expensive outcome.

PART V

Implications and the cXc view

By segment, a short disclosed view from cXc, and a glossary.

18. What this means by segment

The Client money paper's market segmentation — large London Market brokers, UK regional retail brokers, and specialty and MGA-structured intermediaries — applies equally to the platform-consolidation question this paper addresses, with a distinct minimum viable operating model for each.

Large London Market brokers

- **Priority —** multi-currency, multi-syndicate bordereaux reconciliation integrated with Lloyd's Blueprint Two infrastructure as it matures, not retrofitted once the standard settles.
- **Highest-risk layer —** the counterparty layer, given the volume of coverholder and binder relationships a large broker typically manages, and the corresponding TOBA and DUA sprawl documented in Section 6.
- **Realistic destination —** dedicated finance and compliance functions with platform-native tooling, not spreadsheet-augmented versions of an AMS built for a simpler distribution model.

UK regional retail brokers

- **Priority —** commission reconciliation and Consumer Duty MI, given retail brokers' direct exposure to the fair-value evidence standard Section 8 describes.
- **Highest-risk layer —** finance, specifically the manual spreadsheet-based commission matching pattern documented in Section 4, which recurs at nearly every scale in this segment.
- **Realistic destination —** a platform, or a well-integrated point-solution combination, that produces continuous MI as a by-product of normal operation — Scenario B from Section 16 is often the realistic near-term destination, with Scenario C as the medium-term goal.

Specialty and MGA-structured intermediaries

- **Priority —** the counterparty and evidence layers together, given this segment's specific exposure to the FCA's Q2 2026 delegated authority and remuneration review, covered in the CASS paper's Section 9.
- **Highest-risk layer —** authority-as-code architecture (see the MGA companion paper) sitting on top of, or integrated with, the finance and counterparty layers this paper describes — the two are complementary, not competing, investments.
- **Realistic destination —** a platform built around the delegated authority relationship from the outset — this is the segment where Scenario C's full unification argument is strongest, because the segment's growth rate makes Scenario A's compounding cost accumulate fastest.

The common thread across all three segments

Every segment's minimum viable operating model converges on the same underlying requirement: structured data behind commission, counterparty, and evidence workflows, rather than a spreadsheet standing in for a system. The segments differ in which layer is most urgent and how much of the capability map is already addressed by existing tooling — not in whether the underlying platform-consolidation argument applies.

19. A short note on the cXc view

cXc — a trading name of conXchange Pvt Ltd — operates at the distribution platform layer this paper has spent most of its pages examining. This closing section is deliberately brief and explicitly disclosed as commercially interested: readers should treat the propositions below with the same scrutiny they would apply to any vendor's self-assessment, and seek independent validation before acting on them.

Where we think the argument in this paper leads

- **Unification beats integration —** the capability map in Section 14 is not a feature checklist to be satisfied piecemeal; it is a description of what a single system of record needs to do, and the firms that treat it as an integration exercise across several point solutions tend to rebuild the same fragmentation one layer up.
- **The regulatory, operational, and AI arguments are one argument —** the CASS paper's evidence, this paper's spreadsheet-error research, and the AI paper's data-quality precondition are three readings of the same underlying gap. A platform decision made to satisfy only one of these arguments will likely need to be revisited when the next one becomes urgent.
- **Sequencing by damage, not by ease, is the harder but more valuable discipline —** firms that sequence their platform consolidation by which layer is easiest to migrate consistently under-address the layers — finance, counterparty, evidence — actually producing the cost this paper documents.

Where to engage

Readers weighing their own platform consolidation decision — whether that means evaluating cXc, one of the incumbents mapped in Section 15, or an internal build — can reach cXc at cxc.services. The research team responds personally and is glad to discuss the specific capability gaps a firm's current operating model presents, independent of whether that conversation leads to a commercial engagement.

The paper in one line

The spreadsheet was the right tool for a distribution business at a smaller scale, an earlier stage of regulatory scrutiny, and a lower bar for AI-readiness. None of those three conditions hold for most of this market today, and the evidence in this paper — three decades of spreadsheet-error research, this market's own audit findings, and the regulatory tailwinds converging in 2026 — says so independently, from three different directions.

20. Glossary of working terms

Defined at first substantive use throughout this paper; collected here for reference.

Term	Working definition
Bordereau(x)	A periodic report summarising premium, exposure, or claims data exchanged between an MGA or coverholder and a carrier partner.
TOBA	Terms of Business Agreement — the contractual document setting out terms between an intermediary and its counterparties.
DUA	Delegated Underwriting Authority Agreement — the contract granting an MGA or coverholder authority to bind business on a carrier's behalf.
CASS 15	The FCA's safeguarding regime for payment and e-money institutions, in force from 7 May 2026. See the companion paper Client money in a digital-first world for full treatment.
Blueprint Two	Lloyd's multi-year programme to standardise coverholder onboarding, bordereaux processing, and market settlement infrastructure.
CER (Cell Error Rate)	The percentage of spreadsheet cells found to contain an error, in the Panko spreadsheet-error research literature.
SER (Spreadsheet Error Rate)	The percentage of spreadsheets found to contain at least one error, in the same research literature.
MGA / MGU	Managing General Agent / Managing General Underwriter — an intermediary holding delegated underwriting authority from a carrier.
PSP / EMI	Payment Service Provider / Electronic Money Institution — the regulated entities handling premium collection on an intermediary's behalf.
AMS	Agency Management System — the core distribution-layer software platform (Applied Epic, Vertafore AMS360, Acturis, and equivalents).
Delegated authority	An arrangement under which a carrier grants an MGA or coverholder the authority to bind risks on its behalf, within defined terms.
Consumer Duty	The FCA's cross-cutting consumer-protection framework, requiring firms to evidence good outcomes across four defined outcome areas.

References

All sources accessed July 2026. Inline citation numbers map to entries below. Numbering follows the order sources were first cited and is not fully sequential where a source was reused from the cited companion papers' own reference lists.

[1] Peakflo, "Insurance MGA Finance Automation: Bordereaux Processing, Premium Flow, and Commission Reconciliation" (2026). https://peakflo.co/blog/insurance-mga-finance-automation

[2] Inaza, "Beyond the Spreadsheet: Modernizing Bordereau Submissions for MGAs and Reinsurers" (2026). https://www.inaza.com/blog/beyond-the-spreadsheet-modernizing-bordereau-submissions-for-mgas-and-reinsurers

[3] Selectsys, "Commission Tracking in Insurance for MGAs, Wholesalers, and Carriers" (2026). https://www.selectsys.com/resources/commission-tracking-in-insurance

[4] Enrollhere, "Insurance Commission Reconciliation: A Practical Guide" (2026). https://enrollhere.com/blogs/insurance-commission-reconciliation-guide

[5] US Tech Automations, "MGAs: How to Automate Carrier Submissions and Binding 2026", citing McKinsey's 2024 Insurance Operations Efficiency Report and the Insurance Information Institute's 2025 Fact Book. https://ustechautomations.com/resources/blog/automate-mgas-carrier-submissions-and-binding-2026

[6] CT Acquisitions, "Insurance Agency and Broker M&A Multiples Report 2026", citing PitchBook and Insurance Journal data on MGA/MGU transaction multiples. https://ctacquisitions.com/guides/insurance-agency-ma-multiples-2026/

[7] FTI Consulting, "Challenges Outsourcing to Managing Agents" (2026). https://www.fticonsulting.com/insights/articles/challenges-risks-outsourcing-managing-general-agents

[8] R. Panko, "What We Know About Spreadsheet Errors", summarising field-audit research 1995-2004 (cell error rate 5.2%, spreadsheet error rate 94%). https://www.researchgate.net/profile/Raymond-Panko/publication/228662532_What_We_Know_About_Spreadsheet_Errors

[9] R. Panko, "What We Don't Know About Spreadsheet Errors Today: The Facts, Why We Don't Believe Them, and What We Need to Do" (arXiv, 2016). https://arxiv.org/pdf/1602.02601

[10] "Error Estimation in Large Spreadsheets using Bayesian Statistics" (arXiv), summarising Panko 2005 and Powell 2007 field-audit findings. https://arxiv.org/pdf/0908.1190

[11] D. R. Oliver, "Spreadsheets: The Hidden Risk in Your Business" (Medium, 2025), on Panko's overconfidence research and the JPMorgan Chase London Whale incident. https://medium.com/@davidroliver/spreadsheets-the-hidden-risk-in-your-business-989a11af99f2

[12] L. Pryor, "Spreadsheet error rates" (2003), summarising aggregate Panko-cited research. https://www.louisepryor.com/2003/03/24/errorrates/

[13] Global Growth Insights, "Life & Health Insurance Agency Management Software Market 2026". https://www.globalgrowthinsights.com/market-reports/life-health-insurance-agency-management-software-market-103343

[14] Business Research Insights, "Insurance Agency Management Software Size, Share & Analysis 2035 Report". https://www.businessresearchinsights.com/market-reports/insurance-agency-management-s-market-105144

[15] Verified Market Reports, "Insurance Agency Management Systems Market Size & Forecast". https://www.verifiedmarketreports.com/product/insurance-agency-management-systems-market/

[16] DataIntelo, "Global Insurance Agency Management Software Market", on European regulatory drivers and Lloyd's Blueprint Two. https://dataintelo.com/report/global-insurance-agency-management-software-market

[17] US Tech Automations, "Best Insurance Agency Management Software: Save 25% Time 2026", on Applied Epic, Vertafore AMS360, and implementation cost benchmarks. https://ustechautomations.com/resources/blog/best-insurance-agency-management-software-2026

[18] Turris, "Insurance Distribution Management Software Providers 2026". https://www.turrisfi.com/insurance-distribution-management-software-providers-2026

[19] Acturis, company website, on UK premium volume, partner count, and user base. https://www.acturis.com/

[20] 6sense, "Acturis - Market Share, Competitor Insights in Insurance Administration And Management". https://6sense.com/tech/insurance-administration-and-management/acturis-market-share

[21] Regure, "Best Insurance Policy Administration Software 2026", comparing Guidewire, Acturis, Open GI, Applied Epic, Sapiens, Duck Creek, and Insurity. https://www.getregure.com/blog/best-insurance-policy-administration-software-2026/

[22] Astorg, "Astorg Increases Investment in Acturis alongside its Founders and Management" (2026). https://www.astorg.com/news/astorg-increases-investment-in-acturis-alongside-its-founders-and-management

[23] cXc Research, "Client money in a digital-first world" (2026), the companion paper on CASS 5, CASS 7, and the CASS 15 safeguarding regime. https://www.cxc.services/research

[24] cXc Research, "Practical applications of ML in insurance distribution" (2026), the companion paper on AI applications and the integration-tax pattern. https://www.cxc.services/research

[25] cXc Research, "AI in the delegated authority stack" (2026), the MGA companion paper. https://www.cxc.services/research

[26] SubscriptionFlow, "How to Choose MGA Insurance Software For Multi-Carrier Programs" (2026). https://www.subscriptionflow.com/2026/06/how-to-choose-mga-insurance-software/

[38] Advanced Professional Accounting Services, "Bookkeeping for Insurance Agencies", on commission reconciliation practice. https://www.advancedprofessional.net/bookkeeping-insurance-agencies.htm

[42] SubscriptionFlow, "How to Choose MGA Insurance Software For Multi-Career Programs", on inflexible commission setup in standard policy administration systems. https://www.subscriptionflow.com/2026/06/how-to-choose-mga-insurance-software/

[81] FCA, "Consumer Duty board reports: good practice and areas for improvement" (2026). https://www.fca.org.uk/publications/good-and-poor-practice/consumer-duty-board-reports-good-practice-areas-improvement

[82] FCA, "Delivering good outcomes for customers in vulnerable circumstances - good practice and areas for improvement" (2025). https://www.fca.org.uk/publications/good-and-poor-practice/delivering-vulnerable-customers

The distribution series

This paper is part of the cXc Research distribution series. Companions include Client money in a digital-first world (CASS 5/7/15), Practical applications of ML in insurance distribution (the AI parent paper), AI in the delegated authority stack (the MGA companion), and AI in distribution: the technology architecture (the technical architecture companion). All available at cxc.services/research.

cXc

conXchange Pvt Ltd · cxc.services