

Why agentic AI turns spreadsheet chaos into an urgent governance problem, and how to solve it with intelligent software, fusion teams, and a cadence that actually holds.

Managing Director, Digital Insights & Transformation Ltd

Executive Summary

From spreadsheets to systems of intelligence — Why agentic AI turns spreadsheet chaos into an urgent governance problem, and how to solve it with intelligent software, fusion teams, and a cadence that actually holds

The problem

In every regulated organisation I work with, a material share of critical decisions still flows through spreadsheets that no one truly owns, reviews, or governs.


Wholesale platform replacement is too slow. Bans do not survive contact with reality. Custom build does not scale. Agentic AI raises the stakes again: an agent does not notice when the data beneath it is wrong, it simply acts on it, at speed and at scale.

The interesting question is not 'how many spreadsheets do we have?' It is 'which decisions, human or agent-driven, are we making, on what evidence, and through which controls?'

My thesis — treat spreadsheet risk as a governance issue, not an IT issue

If a workbook materially shapes a decision the board is accountable for, then the board is accountable for the workbook. CPS 230 in Australia, BS11 in New Zealand and PS6/21 in the UK now expect exactly this discipline. Most workbooks cannot withstand the scrutiny.

The three-part answer



- **Intelligent software** — overlay technology (e.g. Clear Dynamics, Anaplan) that converts the high-risk workbook into a governed application without disturbing systems of record.
- **Fusion teams** — small, mixed-discipline units chartered to hunt and convert the highest-risk shadow spreadsheets in a defined decision domain.
- **A governance cadence** — a visible rhythm at ExCo (monthly), operational risk committee (quarterly), audit & risk committee (half-yearly), and board (annually) on a single inventory.

Six actions a single executive can sponsor this quarter

- **1. Reframe ownership** into operational risk and resilience (CRO-led, CFO/COO supporting).
- **2. Build a decision-led inventory** of the top-quartile workbooks, capped at 50 entries to force prioritisation.
- **3. Stand up one fusion team** against a coherent decision domain, resourced for at least two quarters.
- **4. Adopt intelligent software for conversions.** Select against the five governance properties, not the vendor brand, and confirm the platform extends the same governance to any AI agent that later consumes the converted application.
- **5. Install the four-tier cadence** with the same dashboard at every level.
- **6. Connect to operational resilience reporting** (CPS 230-equivalent) from day one — avoid parallel programmes.

If your board is asking these questions, get in touch

I work with boards and executive teams across Australia, New Zealand and the United Kingdom on exactly this problem — independent, technology-agnostic, and sized to fit the question. The most useful first conversation is usually short and confidential: which decisions in your organisation feel least defensible today, and what a credible path forward might look like.

Marco Ciobo · marco@marcociobo.com · marcociobo.com · Digital Insights & Transformation Ltd

Key findings

In nearly every large, regulated organisation I work with, I find the same quiet contradiction at the core. The organisation has invested heavily in core platforms, data warehouses, and increasingly AI capability, yet a material share of its critical business decisions still flows through spreadsheets that no one truly owns, reviews, or governs.

I do not think this is any longer just an operational nuisance. As automation and AI accelerate, the gap between how decisions are made and how they are seen to be made widens. Regulators, auditors, and boards are starting to ask sharper questions about model lineage, data provenance, and accountability. The honest answer in many institutions is that critical numbers still originate in a workbook on someone's laptop.

My argument in this paper rests on four key points.

- Spreadsheet sprawl is not an IT hygiene problem. It is a governance problem with strategic, regulatory, and reputational consequences.
- Wholesale platform replacement is the wrong response. It is slow, expensive, and rarely closes the gap that actually matters.
- I treat spreadsheets that drive critical decisions as high-risk by default. The pragmatic way forward is to convert them into governed, auditable, and scalable applications using an intelligent overlay, without disrupting existing systems of record.
- Technology alone is not the answer. The organisations I see succeed combine three things: intelligent software, fusion teams that hunt and convert the highest-risk workbooks, and a governance cadence that holds the executive committee and the board to a visible measure of progress.

I have written this paper for boards, CEOs, CFOs, COOs, CROs, and CIOs of mid-to-large organisations operating in regulated environments.

My practice, Digital Insights & Transformation Ltd, is deliberately vendor-agnostic, entirely independent, and strategy- and governance-focused. Where I cite specific firms tackling this space, I do so only to illustrate that others have recognised the problem and are building viable solutions. My aim is to set out operating-model and governance recommendations that stand on their own and are technology-agnostic.

1. The quiet contradiction at the heart of regulated organisations

When I walk into the finance, risk, treasury, capital, pricing, regulatory reporting, or operations function of a typical bank, insurer, utility, or infrastructure operator, I tend to find the same picture. A handful of trusted analysts are running spreadsheets that drive numbers presented to the board, used in regulatory submissions, or fed into downstream automation. The workbooks have grown organically over the years. They contain undocumented logic, manual overrides, hard-coded assumptions, and external feeds glued together in the hope that they work, largely by trial and error rather than any systematic development methodology. When they break, they tend to be fixed manually, and only when the failure becomes visible.

These tools deliver real value. They are fast, flexible, and shaped by domain experts who understand the question better than any platform team. That is precisely why they have proved so resilient. The problem is not the spreadsheet. The problem is the absence of the controls that would normally surround a system and a record of this importance.

In my work, the pattern usually has six characteristics:

- Critical-path workbooks are owned by individuals, not by accountable functions.
- Logic is embedded in formulas and macros that are rarely, if ever, peer-reviewed.
- Inputs come from a mix of system extracts, email attachments, and manual entry.
- Versions proliferate among team members and across departments, between developers and ultimate consumers, with no canonical or often traceable source of truth.
- Approvals happen verbally or by email, not as part of an auditable workflow.
- When the owner is on leave or leaves the organisation, the function becomes brittle, often fractures, and critical data and insights are lost.

My experience is that when you layer automation, robotic process automation, and now generative AI on top of this foundation, the risk does not stay constant, it compounds.

These technologies are powerful, but atop a shaky data foundation the risk proliferates and is often amplified. Every new step that consumes spreadsheet output without challenging it inherits its assumptions, its errors, and its silence.

The arrival of agentic AI sharpens this risk further. Where a human analyst consuming a spreadsheet might notice something is wrong, an agent does not. It executes, at speed and at scale, against whatever data it is given. If that data originates in an ungoverned workbook with undocumented logic, uncontrolled overrides, and no audit trail, the agent inherits every one of those weaknesses and acts on them. Boards approving the use of AI agents in regulated processes need to ask a prior question: is the data those agents will act on governed? In most organisations I work with, the honest answer is not yet.

Case study — UK Public Health England COVID-19 Excel incident (October 2020)

During the early second wave of the pandemic, Public Health England (PHE) was importing commercial laboratory test results into England's central reporting system used for national COVID statistics, dashboards, and NHS Test and Trace contact tracing operations. One of the ingestion pipelines relied on the legacy Excel .xls format, which is capped at 65,536 rows per worksheet. When daily inbound data exceeded the cap, rows beyond the limit were silently and without anyone noticing for some time, truncated.

Impact

The net effect was approximately 15,841 positive test cases recorded between 25 September and 2 October 2020 were omitted from daily reporting. Infected individuals were not promptly referred to contact tracers, and local authorities operated on incomplete outbreak data during a critical growth phase. Subsequent analysis suggested the delay may have contributed to tens of thousands of additional infections and over a thousand additional deaths, although PHE disputed the strength of those causal claims.

Why I include it

In my view, the deeper cause was not Excel itself; it was the architectural choice to place a spreadsheet inside a critical national data pipeline without adequate validation, monitoring, or escalation controls.

The incident has become a standard reference for data governance, operational resilience, and the dangers of using end-user computing tools for enterprise-scale workflows. It is exactly the pattern I argue regulated organisations should pre-emptively eliminate.

2. Treat spreadsheet risk as a governance issue, not an IT issue

If a workbook materially shapes a decision the board is accountable for, then the board is accountable for the workbook. Treating it as an IT hygiene problem under-serves the question.

I have seen most organisations attempt a spreadsheet inventory at some point. The output is usually a long list, ranked by file size or user count, that is then handed to a delivery team and quietly deprioritised. In my experience this pattern fails for three reasons.

First, the framing is wrong. The interesting question is not “how many spreadsheets do we have?” It is “which decisions are we making, on which evidence, and through which controls?” That is a question of governance, not of file management.

Second, the wrong people own the response. Treating the issue as an IT clean-up means it competes for budget with platform projects backed by louder sponsors, and almost always loses. The accountable owners, the CFO, COO, Chief Risk Officer, Head of Regulatory, are usually not in the room.

Third, the success measure is wrong. Removing spreadsheets is not the goal. The goal is to know which decisions you can defend, audit, and reproduce. A governed spreadsheet that meets the bar is a perfectly acceptable outcome. An ungoverned application is not.

My argument that spreadsheet risk is a governance problem, not an IT problem, is supported by two decades of high-profile incidents in which fragile end-user workbooks have shaped material decisions.

The two cases I have set out below, JPMorgan’s “London Whale” trading loss and the Reinhart–Rogoff economics paper, are widely taught in business schools around the world, they are well documented, and instructive precisely because both organisations had ample resources to do better at the time.

Case studies — when spreadsheets shaped material decisions

JPMorgan “London Whale” (2012)

JPMorgan’s Chief Investment Office accumulated synthetic credit derivative positions that ultimately produced losses in excess of US\$6 billion.

A US Senate Permanent Subcommittee on Investigations report later identified that the bank’s Value at Risk (VaR) model had been re-implemented in spreadsheets that involved manual copy-paste between workbooks, undocumented formula changes, and a calculation error in which sums had effectively been used in place of averages.

The result was a material understatement of trading risk through a period in which exposures were growing rapidly. The episode became a defining case in financial-services discussions of model risk, spreadsheet risk, and end-user computing governance, and contributed to subsequent industry focus on BCBS 239 risk data aggregation standards.

Reinhart and Rogoff “Growth in a Time of Debt” (2010–2013)

The Reinhart–Rogoff paper argued that economic growth drops sharply when government debt-to-GDP exceeds roughly 90 per cent, a finding widely cited to justify austerity policies in Europe and elsewhere.

In 2013, graduate student Thomas Herndon (UMass Amherst) attempted to replicate the analysis and discovered an Excel cell-range error in which several countries had been omitted from an average calculation, alongside selective data treatments and unusual weighting choices.

Once corrected, the dramatic “cliff” at the 90 per cent threshold largely disappeared. The episode is now taught widely in economics, public policy, and data science as a case in reproducibility, spreadsheet fragility, and the importance of transparent code and data.

The common pattern I see

In all three incidents, PHE, JPMorgan, and Reinhart–Rogoff, I read the failure as not really an “Excel problem.”

Each was a governance failure in which a material decision (a public-health response, a multi-billion-dollar risk position, a global policy argument) was driven by a fragile spreadsheet-based process that lacked validation, reproducibility, and independent challenge. The cost of that fragility was paid in lives, capital, and credibility respectively.

Boards in Australia, New Zealand, and the United Kingdom now have sharpened reasons to reframe this question. APRA’s Prudential Standard CPS 230 (Operational Risk Management) places explicit obligations on regulated entities to identify, manage, and test the operational risk profile of critical operations and material service providers.

Equivalent obligations sit in the Reserve Bank of New Zealand’s BS11 expectations and the UK PRA’s operational resilience regime. From where I sit, spreadsheet-driven critical operations sit squarely inside this perimeter.

My implication is unambiguous: if a spreadsheet supports a critical operation, then the controls around it must withstand the same scrutiny as the controls around any other system that supports a critical operation. Most workbooks I see cannot.

3. Why the obvious responses do not work

Three responses tend to dominate when I sit down with a leadership team facing this problem. Each looks reasonable. In my experience, each on its own fails.

3.1 Replace with a major platform programme

The instinct is to retire the workbooks by embedding their logic into the next ERP, core banking, billing, or actuarial platform.

In practice, the platform roadmap is already crowded, the spreadsheet logic is poorly documented, and the affected teams cannot afford to wait. By the time the platform release arrives, the spreadsheet has evolved further.

Replacement programmes also tend to deliver capability, not control: the data may move, but the decision rights, audit trail, and override governance frequently do not.

3.2 Ban or restrict spreadsheets

Edicts to remove or restrict spreadsheet use do not survive contact with reality.

The work still has to happen, and the workbooks reappear in OneDrive, SharePoint, or shadow team sites. A request from an executive for a rapid sensitivity analysis of a business case being presented to the board results in a workbook that grows in complexity as new assumptions and implications are added after the original request. A ban does not eliminate the original conditions or necessity of creation. Worse, the controls that previously existed in plain sight migrate into less visible places.

Restriction without an alternative is a control regression, not an improvement.

3.3 Build something custom

An in-house build can look attractive when a commercial off-the-shelf solution does not fit. The reality is that the organisation rarely has a unit-economics model that supports building and operating purpose-built applications at the rate the business needs them. The result is a backlog of half-built tools that the platform team cannot sustain.

My observation is that all three responses miss the same point: the right unit of intervention is neither the spreadsheet nor the platform. It is the decision and the governance that surrounds it.

4. The three-part answer: intelligent software, fusion teams, and a governance cadence

In my experience, technology alone, however excellent, does not deliver a durable answer. Nor does a fusion team without the right tools, nor a governance forum without the throughput to act on its decisions. The organisations I see move the needle on spreadsheet risk combine all three:

- **Intelligent software** — a fit-for-purpose technology approach that converts the high-risk workbook into a governed, auditable application without disturbing the underlying systems of record.
- **Fusion teams** — small, durable, mixed-discipline units chartered to hunt and convert the highest-risk shadow spreadsheets in a defined decision domain.

- **A governance cadence** — a visible rhythm at the executive committee, operational risk committee, audit and risk committee, and board, anchored on a single dashboard and a single inventory.




Figure 1. The three-part answer — where the three combine, decisions become defensible.

In the rest of this section I set out each of the three in turn, and then describe the four-stage remediation playbook I use to link them.

4.1 Intelligent software: the overlay that converts the spreadsheet into a governed application

Over the past few years I have watched a more pragmatic alternative emerge to platform replacement and custom build.

It treats the high-risk spreadsheet as the specification and converts it, quickly, and without disturbing the underlying systems of record, into a governed, auditable application. The user-facing experience can remain familiar; what changes is the layer underneath.

I have seen this category described as “intelligent overlay” or “intelligent software development.” The defining property is that the application is generated from a description of the decision and its data, rather than coded by hand from a blank page. That is what makes it economic at the scale most regulated organisations actually need.

In my view, an intelligent overlay has to deliver five capabilities that a spreadsheet cannot:

- **Identity and access.** Every action is attributable to a named user with an authorisation profile that the organisation can defend.

- **Versioning and lineage.** Inputs, calculations, overrides, and outputs are tracked over time, not lost between Save As copies.
- **Workflow and approval.** Human-in-the-loop checkpoints are explicit and auditable rather than verbal or email-based.
- **Data integration.** The application connects to authoritative sources rather than ingesting copy-pasted extracts.
- **Operability and resilience.** The application is monitored, backed up, and supportable through normal IT service management — not dependent on one person's laptop.

Critically, the overlay does not replace the system of record. It sits alongside it. That, in my view, is what makes it implementable in regulated settings where wholesale change is unaffordable in time, money, and risk.

The emergence of credible exemplars

These insights are not mine alone. I have watched a number of firms emerge that focus directly on this problem and are building viable, commercially supported solutions.

The theme has become more urgent with the accelerating effects of machine learning and AI, and the governance risks they pose to boards. I set out two illustrative, though by no means exclusive, exemplars below.

Clear Dynamics

Clear Dynamics is an Australian intelligent software company with a strong heritage in highly regulated financial services. Rather than leading with a platform, Clear Dynamics works alongside customer teams to deliver outcomes.

Small, focused groups work closely with the business to understand the problem, convert the high-risk workbook into a governed application, and then host and maintain it under a SaaS model. The customer always owns their IP; Clear Dynamics keeps the application evergreen.

The governance properties I have set out above are built in by design, not added afterwards as controls. I have seen this model applied at Bank of Queensland for decision-critical workflows, and at Qube Logistics for a complex supply chain tracking process, both converted from spreadsheet-driven operations to governed applications without triggering a multi-year platform programme.

Clear Dynamics has also extended this capability to agentic AI governance, building guardrails into applications that use AI agents, including human-in-the-loop checkpoints at defined workflow stages and controls on token consumption. For boards beginning to approve AI agent use in regulated

processes, this kind of governance infrastructure is not optional. It is the same discipline as spreadsheet governance, applied one layer up.

Anaplan

Anaplan is a US-headquartered enterprise planning platform widely used to replace spreadsheet-driven planning, forecasting, and operational modelling.

Its connected-planning model combines a multi-dimensional calculation engine with role-based access, audit logging, workflow, and integration with authoritative data sources, in other words, the same five governance properties I set out above. I have seen Anaplan deployed at scale in financial services, telecommunications, and government for use cases including workforce planning, financial planning and analysis, supply-chain modelling, and capital allocation.

To me, its main contribution to this category is in showing that connected, governed planning can replace large estates of inter-linked workbooks without disturbing systems of record.

I refer to Clear Dynamics and Anaplan here as established, credible exemplars of the category. Other approaches exist, and I would guide selection by the five governance properties above rather than by vendor brand. What matters for the board is that the chosen technology demonstrably meets those properties in the regulated context the organisation actually operates in.

4.2 Fusion teams: the operating unit that targets shadow spreadsheets

I have rarely seen spreadsheet sprawl solved by a project team in IT. It is solved by a small, durable, mixed-discipline unit with the authority to redesign the decision and the engineering capacity to implement it. The pattern I see travel best is the fusion team, five to nine people, accountable to a business owner, with engineering, data, and risk skills embedded, chartered explicitly to target the shadow spreadsheets that sit behind the organisation's highest-risk decisions.

In my experience, a fusion team that works has the following composition:

- **Business owner (accountable).** A senior leader from the function that owns the decision (e.g. Head of Treasury, Head of Capital, Head of Pricing). They sign off on scope, acceptance, and go-live, and are accountable for the controls afterwards.
- **Process owner / domain expert.** Often the person who currently updates the spreadsheet. Their tacit knowledge is the asset; the team protects it and codifies it.
- **Engineer or product builder.** Capable of implementing the overlay quickly, working from the spreadsheet as a specification.

- **Data engineer.** Owns the connections to authoritative sources, replacing manual extracts and email attachments.
- **Risk and controls partner.** Translates the workbook's implicit controls into explicit, testable controls in the application.
- **Governance / audit liaison (part-time).** Ensures the team's artefacts are the ones internal audit and regulators will later request.

Three operating principles, in my view, determine whether a fusion team succeeds or quietly disbands:

- **Mission, not project.** The team is chartered against a portfolio of related decisions (for example, “regulatory liquidity reporting” or “pricing exceptions”) rather than a single workbook. Continuity matters more than novelty.
- **Single-threaded ownership.** One named business owner, one named technical lead. Multi-headed sponsorship is the most reliable way to lose momentum.
- **Embedded, not seconded.** Members are co-located (physically or virtually) for the duration of the mission, not occasional contributors squeezed between other duties.

A useful diagnostic question for the executive committee: for each of our top ten spreadsheet-driven decisions, can we name the fusion team that owns the path to a governed alternative? If not, the issue is organisational, not technical.

4.3 The shadow-spreadsheet remediation playbook

The playbook I use works through the portfolio in four sequential stages. Each stage has clear inputs, outputs, and exit criteria.

I have applied the same shape to a single workbook over days and to a portfolio over quarters; the shape itself does not change. This is the engine that connects the technology in 4.1 to the governance cadence in 4.4.

Stage 1 — Discover and triage

I build a single, executive-visible inventory of decision-critical workbooks. I do not try to catalogue every spreadsheet in the organisation; that is the trap that has defeated most prior efforts I have reviewed. I start from the decisions, not the files.

For each candidate I capture:

- the decision it supports and the named accountable executive;
- the regulatory, financial, or reputational consequence of error;
- the volume and frequency of use;
- the current control environment (or absence thereof);
- the current owner and the “bus factor” if that owner is unavailable.

I then triage the inventory into a simple severity-by-likelihood grid. The output of Stage 1 is a prioritised remediation queue that the executive committee has seen, debated, and signed.

Stage 2 — Stabilise the highest-risk workbooks in flight

In my work, some workbooks cannot wait for full remediation. The point of Stage 2 is to apply control hygiene so the residual risk is tolerable while the queue is worked through. This is not a substitute for remediation; it is a holding action.

Practical interventions I use include:

- named primary and secondary owners, recorded in the inventory;
- version locking and change logs;
- documented assumptions, with sign-off by the accountable executive;
- second-pair-of-eyes review for each material change;
- pre-defined override rules, captured rather than negotiated each cycle.

Stage 3 — Convert to a governed application using intelligent software

For the workbooks that warrant it, I have the fusion team convert the spreadsheet into a governed application using the intelligent software approach in 4.1. The spreadsheet becomes the specification; the application becomes the system of record for that decision. The work is iterative, with the business owner accepting each increment.

In my view, a converted application should at a minimum demonstrate:

- named users, role-specific access, and full action logging;
- live connections to authoritative data sources, replacing manual extracts;
- explicit workflow with peer review and approval steps;
- controlled override paths with reason codes and audit trail;
- reproducibility — any output can be recomputed from preserved inputs at a point in time;
- integration with the organisation's standard incident, change, and access management.

Crucially, I only retire or restrict the source spreadsheet once the application has been operating in parallel long enough to satisfy the business owner and risk partner. Cutover, not surprise, is the rule.

Stage 4 — Embed and sustain

Stage 4 is where I see most remediation programmes quietly fail. The application is delivered, the team disperses, and within twelve months a new generation of side-spreadsheets has grown around the edges. Sustaining the gain, in my experience, requires four things:

- the application is on the standard operating runbook (incident, change, monitoring, backup) — not on a side path;
- the business owner has continuing accountability for the control environment, not just for the build;
- spreadsheet creation around the application is monitored, with new candidates folded back into the inventory;
- the fusion team's next mission is queued, so the capability does not evaporate.

4.4 Governance cadence and accountability

In my experience, the single biggest determinant of sustained progress is the cadence at which the executive committee and the board look at the inventory. The programmes I have seen succeed share the following rhythm:

Forum	Cadence	What is reviewed
Executive committee	Monthly	Top-quartile inventory items, in-flight remediations, RAG status, blockers requiring executive intervention.

Forum	Cadence	What is reviewed
Operational risk committee	Quarterly	Risk-weighted heatmap, control environment changes, incidents traced back to workbook root cause.
Audit / risk committee of the board	Half-yearly	Trend in inventory size and severity, remediation throughput, regulatory engagement summary.
Board	Annually (with breach triggers)	Strategic posture, residual risk appetite, capital and resource implications.

A few principles, in my view, make this cadence productive rather than performative:

- **Same view, every time.** The inventory and dashboard are stable. Executives stop arguing about the chart and start arguing about the substance.
- **Owner present, in person.** Each red item is presented by its accountable executive, not delegated to a programme manager.
- **Evidence-based RAG.** Status changes are tied to defined exit criteria from the playbook, not to opinion.
- **Linkage to operational resilience reporting.** The inventory feeds the organisation's broader operational resilience and CPS 230-equivalent reporting, so it is not parallel work.

Three numbers belong on the dashboard from day one: the count of decision-critical workbooks in scope, the count remediated to a governed alternative, and the count of incidents in the last quarter traced to a workbook root cause. Everything else is supporting detail.

5. What good looks like in twelve and twenty-four months

An organisation that has taken my three-part answer seriously — intelligent software, fusion teams, and a governance cadence — will, by the end of year one, be able to answer four questions cleanly:

- Which decisions in our organisation are spreadsheet-dependent today, and who is accountable for each?

- Which of those have been converted to a governed alternative using intelligent software, and which are stabilised but not yet converted?
- What is our quarterly throughput of conversions, and what is the bottleneck — technology, fusion-team capacity, or executive sponsorship?
- Where, if anywhere, are workbooks still creating audit, regulatory, or reputational exposure?

By the end of year two, I would expect the conversation to have shifted. The question is no longer “how many spreadsheets do we have?” It is “are we comfortable with the residual risk in the workbooks we have chosen to keep?” In my view, that is a defensible governance posture.

6. Recommended actions for boards and executive teams

The recommendations I set out below are deliberately concrete. Any one of them can be sponsored by a single executive in the next quarter.

1. **Reframe ownership.** Move the spreadsheet question out of IT-hygiene and into the operational risk and resilience agenda. My recommended accountable owner is the Chief Risk Officer (or equivalent), supported by the CFO and COO.
2. **Build the inventory from decisions, not files.** Commission a decision-led inventory of the top-quartile workbooks across finance, risk, treasury, capital, pricing, regulatory reporting, and operations. I cap the initial scope at fifty entries to force prioritisation.
3. **Stand up one fusion team.** Charter a single fusion team against a coherent decision domain. Resource it for at least two quarters. I use it to validate the playbook and to build internal credibility.
4. **Adopt intelligent software for conversions.** Avoid a fresh in-house build and avoid waiting for the next platform release. Select a fit-for-purpose intelligent software capability; Clear Dynamics and Anaplan are two credible examples I work with, that clearly meets the five governance properties I have set out in Section 4.1.
5. **Install the cadence.** Stand up the four-tier governance cadence (executive committee, operational risk committee, audit/risk committee, board) and use the same dashboard at every level.
6. **Connect to operational resilience reporting.** Make the inventory and remediation status part of CPS 230-equivalent reporting from day one. Avoid a parallel programme that competes for executive attention.

7. Closing

In my view, spreadsheet chaos is one of the rare problems where the technical fix is now the easier half. Intelligent software has made converting high-risk workbooks into governed applications faster, cheaper, and more reliable than even five years ago. The harder, more valuable half is the operating-model and governance change that surrounds it — the fusion teams that target the shadow spreadsheets, and the cadence that holds the executive committee and the board accountable for visible progress.

Boards that take that harder half seriously will, in my experience, end up with something more durable than a tidier file system. They end up with decisions they can defend.

That is the prize. I believe it is achievable, it is affordable, and — for organisations operating under CPS 230 and its international equivalents — it is increasingly non-negotiable.

If your board is asking these questions, let's talk

If the picture I have set out in this paper looks familiar, spreadsheets behind decisions you cannot fully defend, an operational-resilience programme that does not yet include them, or a board that is starting to ask sharper questions, there are practical next steps you can take inside a single quarter.

Through Digital Insights & Transformation Ltd, I work with boards and executive teams across Australia, New Zealand, and the United Kingdom to do four things: frame the issue for the board, run a fast decision-led inventory of the top-quartile workbooks, stand up the first fusion team, and install the governance cadence that holds the programme to visible progress. My work is deliberately independent and technology-agnostic, and I size it to fit the question, from a single board briefing through to a multi-quarter remediation engagement.

If your board is asking these questions, get in touch. The most useful first conversation I have is usually a short, confidential one about which decisions in your organisation feel least defensible today — and what a credible path forward might look like.

Where the engagement calls for a technology solution, I work with a small number of established intelligent software providers, including Clear Dynamics, that I have seen deliver the five governance properties in Section 4.1 in regulated environments. My advisory work and their delivery capability are separate and complementary.

Marco Ciobo · marco@marcociobo.com · marcociobo.com · Digital Insights & Transformation Ltd

References and further reading

UK Public Health England COVID-19 reporting incident, October 2020

- BBC News (5 October 2020). “Covid: 16,000 coronavirus cases missed in daily figures after IT error.” [bbc.co.uk/news](https://www.bbc.co.uk/news)
- National Audit Office (2021). The government’s approach to test and trace in England — interim report. [nao.org.uk](https://www.nao.org.uk)
- Kelion, L. (5 October 2020). “Excel: Why using Microsoft’s tool caused Covid-19 results to be lost.” BBC News, [bbc.co.uk/news/technology-54423988](https://www.bbc.co.uk/news/technology-54423988)

JPMorgan Chase “London Whale” trading losses, 2012

- US Senate Permanent Subcommittee on Investigations (15 March 2013). JPMorgan Chase Whale Trades: A Case History of Derivatives Risks and Abuses. [hsgac.senate.gov](https://www.hsgac.senate.gov)
- JPMorgan Chase & Co. Management Task Force (16 January 2013). Report of JPMorgan Chase & Co. Management Task Force regarding 2012 CIO Losses. [jpmorganchase.com](https://www.jpmorganchase.com)
- US Government Accountability Office (2014). Lessons from the London Whale (selected discussion). [gao.gov](https://www.gao.gov)

Reinhart and Rogoff “Growth in a Time of Debt”, 2010–2013

- Reinhart, C. M., & Rogoff, K. S. (2010). “Growth in a Time of Debt.” *American Economic Review*, 100(2), 573–578.
- Herndon, T., Ash, M., & Pollin, R. (April 2013). “Does High Public Debt Consistently Stifle Economic Growth? A Critique of Reinhart and Rogoff.” PERI Working Paper 322, University of Massachusetts Amherst.
- Krugman, P. (2013). “The Excel depression.” *The New York Times* (commentary).

Regulatory standards

- Australian Prudential Regulation Authority (APRA), Prudential Standard CPS 230 — Operational Risk Management. [apra.gov.au](https://www.apra.gov.au)
- Reserve Bank of New Zealand, BS11 Outsourcing Policy. [rbnz.govt.nz](https://www.rbnz.govt.nz)
- UK Prudential Regulation Authority, PS6/21 Operational Resilience. [bankofengland.co.uk](https://www.bankofengland.co.uk)
- Basel Committee on Banking Supervision, BCBS 239 — Principles for effective risk data aggregation and risk reporting. [bis.org](https://www.bis.org)

Spreadsheet risk research

- European Spreadsheet Risks Interest Group (EuSpRIG) — Horror Stories archive and research papers. [eusprig.org](https://www.eusprig.org)
- Panko, R. R. (2008). “Spreadsheet Errors: What We Know. What We Think We Can Do.” Proceedings of EuSpRIG.

About this paper

This is an independent positioning piece I have written, drawing on my direct experience leading transformation, AI, and governance work with banks, insurers, utilities, and infrastructure operators across Australia and New Zealand. I write in my capacity as Managing Director of Digital Insights & Transformation Ltd, and as a Senior Advisor to the Bain Advisor Network. The views expressed are entirely my own.

Contact: marco@marcociobo.com | Digital Insights & Transformation Ltd