2026 EDITION

AI-Ready Finance

The Practitioner's Guide for UK and European Mid-Market Finance Teams

Maebh Collins FCA

Fellow Chartered Accountant · Big 4 Trained · Finance Transformation Specialist

EXECUTIVE SUMMARY

AI adoption in finance is accelerating. Whether it is creating value is a different question.

59% of finance functions now use AI in some form, up from 37% in 2023, according to Gartner research. The pressure to act is real. Boards are asking. Vendors are demonstrating. Budgets are being allocated. And yet: 70 to 85% of AI projects fail to meet their objectives, and 95% of generative AI pilots deliver zero measurable financial return. Only 48% make it into production.

59%	~85%	48%
of finance functions now use AI in some form (Gartner, 2025)	of AI projects fail to deliver intended value (Gartner, 2024)	of AI pilots make it into production

The problem is not the technology. The problem is what sits underneath it.

Finance functions that attempt to deploy AI onto fragmented data, undocumented processes, and inconsistent chart of accounts structures do not get better results: they get faster errors. The gap between AI ambition and AI readiness is where most mid-market finance teams are currently sitting, and where most of the money is being wasted.

This guide addresses that gap.

Written for CFOs and Finance Directors in UK and European mid-market businesses. Vendor-neutral. Written by someone who has built and automated finance functions with her own capital at stake, not from the advisory side.

THE CENTRAL INSIGHT

The most counterintuitive finding in this guide: the foundations that determine whether AI delivers value in your finance function are the same foundations that determine whether your management accounts are reliable. A consistent chart of accounts. Documented processes. Clean data. Controls that hold. This is not AI preparation: it is good accounting practice. The sequencing is what most organisations get wrong.

The goal is not to make the case for AI. That case is already made. The goal is to help you answer the question that matters more: is your finance function ready for it?

I

CHAPTER 1: WHERE FINANCE REALLY STANDS WITH AI

The headline figures on AI adoption in finance sound impressive. Dig one level deeper and a more complicated picture emerges.

According to Gartner, 59% of finance functions are now using AI in some form. That number doubled between 2023 and 2024, rising from 37% to 58% in a single year. It is the kind of growth that generates conference keynotes and board-level anxiety. The question worth asking is not how many finance functions are using AI. It is what they are doing with it, and whether it is working.

McKinsey's data makes the gap concrete. 88% of organisations report using AI somewhere in their operations. Only 6% qualify as what McKinsey calls "AI high performers": organisations seeing significant, measurable impact from their AI investments. The gap between using AI and getting value from it is where most finance functions currently sit.

AI ADOPTION IN FINANCE FUNCTIONS: THE PLATEAU

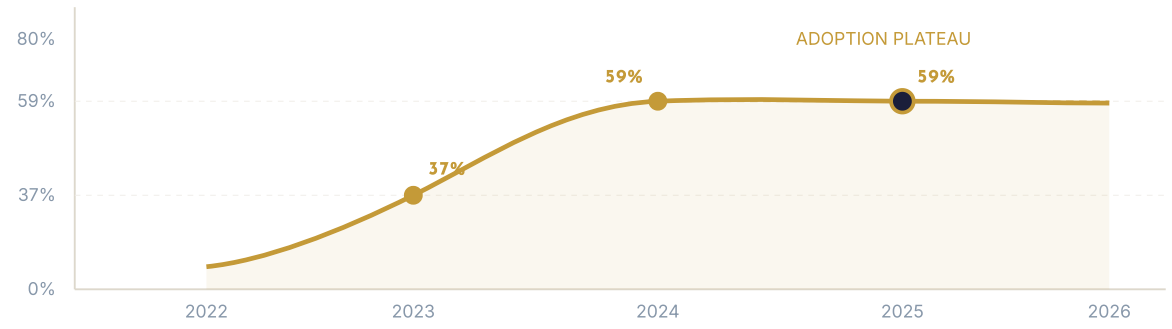

Source: Gartner (2024-2025). Adoption doubled 2023-2024, then stalled. The plateau is diagnostic.

THE PLATEAU IS TELLING YOU SOMETHING

Gartner's most recent data shows adoption has plateaued at 59%. After the sharp acceleration of 2023 and 2024, growth has stalled. This is not unusual in technology adoption cycles. It typically signals one of two things: either the easy use cases have been captured and the harder ones are proving more difficult, or adoption has run ahead of the foundations required to make it work. In finance, both are true.

The organisations that moved earliest on AI in finance are now two to three years into working out what actually delivers results and what does not. They have learned which use cases require clean data and which can tolerate some inconsistency. They have built the governance frameworks their auditors are now asking about. They have changed processes, not just added tools onto existing broken ones. That is a meaningful head start.

For mid-market finance functions that have not yet moved, this creates a specific challenge. The window to learn cheaply, through experimentation with lower stakes, has narrowed. But it also creates an opportunity: the early mistakes have been made and documented by others. You can see what failed and why before you commit budget.

WHERE FINANCE FUNCTIONS ARE ACTUALLY USING AI

Not all AI adoption in finance looks the same. Adoption rates vary significantly by sub-function, and the pattern reveals something important about where the prerequisites are easiest to meet.

Accounts payable automation sits at the top of the adoption curve. This is not surprising. AP is a high-volume, rules-based process with relatively structured data inputs: invoices, purchase orders, payment terms. The matching logic is well-defined. The outcomes are measurable. AI can add genuine value here without requiring a pristine data foundation across the whole function.

Autonomous financial close sits at the bottom. Again, not surprising. Month-end close draws on data from across the organisation: multiple systems, multiple entities, manual journal entries, estimates, and accruals that require human judgement. Automating the close requires that all of those upstream inputs are clean, consistent, and documented. For most mid-market finance functions, they are not.

This pattern holds across use cases. AI performs best where data is structured, processes are documented, and the rules are clear. It performs worst, or fails entirely, where the opposite is true.

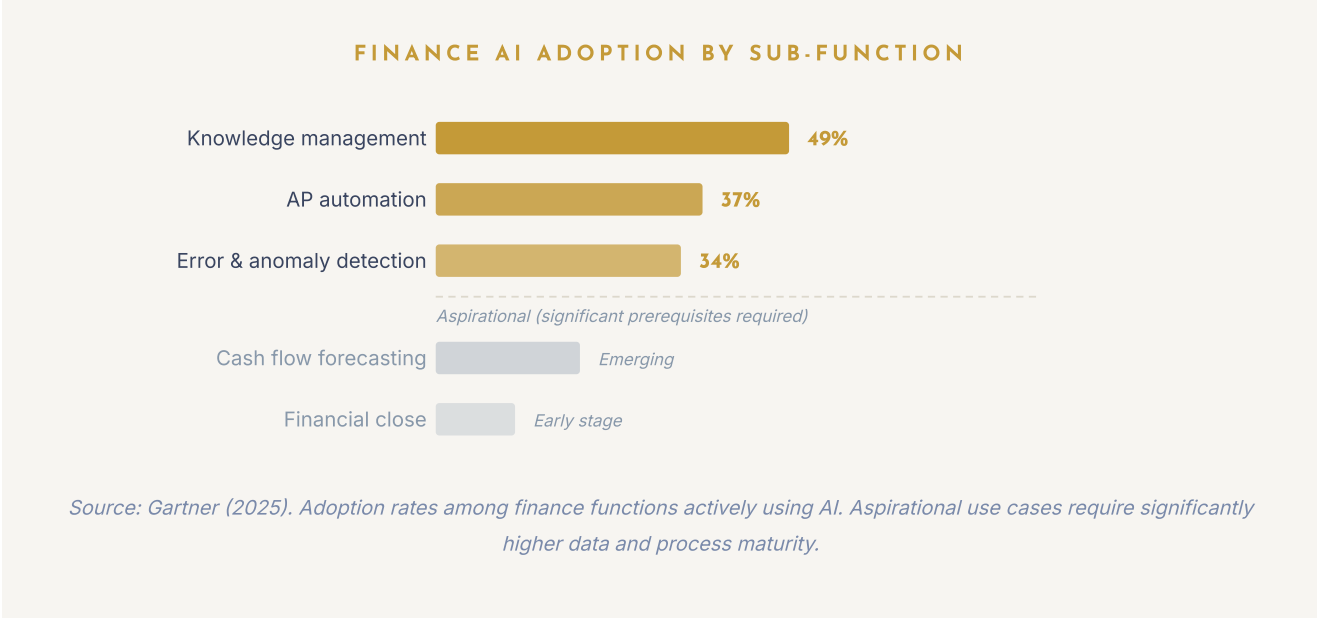

Source: Gartner (2025). Adoption rates among finance functions actively using AI. Aspirational use cases require significantly higher data and process maturity.

THE DATA PROBLEM AT THE CENTRE OF IT

EY's 2025 Tax and Finance Operations Survey asked finance leaders to identify their most significant barrier to AI adoption. 80% named insufficient AI-ready data. Not budget. Not skills. Not technology. Data.

This is consistent with every other major piece of research on AI in finance, and with direct experience of building and transforming finance functions. The data underneath most finance operations is messier than people want to admit: inconsistent chart of accounts structures across entities, supplier records with multiple naming conventions, reconciliation processes that exist in one person's head rather than in documented procedures, ERP configurations that have accumulated years of workarounds.

AI does not clean this up. It exposes it.

The organisations currently struggling with AI pilots in finance are not struggling because the technology does not work. They are struggling because the technology is working exactly as designed, processing the data it is given and producing outputs that reflect the quality of that data.

WHAT THIS MEANS FOR MID-MARKET FINANCE LEADERS

The pressure to adopt AI is real. Boards are asking. Competitors are announcing. Vendors are demonstrating tools that, in a controlled environment with clean sample data, look compelling.

The honest position for most mid-market finance functions in 2026 is this: the foundations are not yet in place. The data is not ready. The processes are not documented. The governance frameworks do not exist. And jumping into AI deployment before addressing those foundations does not accelerate transformation: it accelerates the discovery of problems that were already there.

That is not an argument against AI adoption. It is an argument for sequencing it correctly.

The AI Readiness Self-Assessment in Chapter 7 gives you a structured way to benchmark your function across the five readiness pillars. Before you reach it, the question worth sitting with is not whether to adopt AI: it is whether your finance function is ready for it, and whether the sequence in which you move actually reflects that.

2

CHAPTER 2: THE USE CASES THAT ACTUALLY WORK VERSUS VENDOR HYPE

Every AI vendor in finance can show you an impressive demo. The reconciliation matches in seconds. The anomaly is flagged before the month-end close. The cash flow forecast updates in real time. It looks, in the controlled environment of a sales presentation, like the future of finance.

The demo data is clean. It is structured. It has been prepared specifically to show the tool performing well. Your data is not like that, and the gap between the two is where most AI projects in finance run into trouble.

This chapter covers what the evidence actually shows: which use cases are delivering results, which remain largely aspirational, and what needs to be in place before any of them can work.

THE THREE USE CASES WITH REAL TRACTION

USE CASE	ADOPTION	KEY PREREQUISITE	HYPE VS REALITY
Knowledge management	49%	Digitised, accessible documents. No need for pristine chart of accounts.	DELIVERS
AP automation	37%	Consistent supplier master data and clean chart of accounts mapping.	CONDITIONAL
Error & anomaly detection	34%	Three years of clean, consistently categorised historical transaction data.	CONDITIONAL

DELIVERS = works with standard prerequisites. CONDITIONAL = works only if specific data/process foundations are in place.

Gartner tracks AI adoption across finance sub-functions. Three use cases have reached adoption by more than a third of finance functions currently using AI, making them the most reliably evidenced starting points.

Knowledge management (49% adoption). This covers using AI to search, retrieve, and summarise information across internal documents: policies, contracts, historical reports, technical guidance. Large language models are well-suited to this task. The prerequisites are relatively low compared to transactional use cases. You need documents that are digitised and accessible, but you do not need a pristine chart of accounts. For finance teams spending significant time searching for information, the productivity gains are real and relatively quick to achieve.

Accounts payable automation (37% adoption). AI-assisted invoice processing, three-way matching, and payment scheduling. This is the most mature use case in finance AI, and the results in well-implemented deployments are measurable: faster processing, fewer exceptions requiring manual review, reduced duplicate payments. The prerequisite here is a consistent supplier master data file and a chart of accounts that maps cleanly to incoming invoice categories.

Without those, the matching rate drops and manual intervention increases. Often to the point where the tool creates more work than it saves.

Error and anomaly detection (34% adoption). AI models trained on historical transaction data to flag outliers: duplicate entries, unusual payment patterns, journal entries posted at atypical times or by atypical users. This is a strong use case for finance because the stakes of missing an error are high and the volume of transactions makes manual review impractical. The prerequisite is sufficient clean historical data to train a meaningful baseline. A finance function with inconsistent data categorisation over the past three years will produce a model with a high false-positive rate.

USE CASES WITH REAL POTENTIAL, AND SIGNIFICANT PREREQUISITES

Cash flow forecasting. AI tools are achieving 90 to 95% accuracy in structured cash flow forecasting environments. That figure is worth interrogating. It applies to businesses with consistent revenue patterns, well-integrated ERP and banking data feeds, and historical data that reflects normal operating conditions. It does not apply universally. A business that went through Brexit restructuring, a pandemic, or a major acquisition in the past five years has historical data that is not representative of future patterns. The model will be trained on noise.

Financial close acceleration. Automating elements of the month-end close: journal entry preparation, intercompany reconciliation, variance commentary. The potential here is significant. The prerequisites are among the most demanding of any AI use case: fully documented close processes, consistent account coding, clean intercompany balances, and journal entry templates that have been standardised across the function. Most mid-market finance teams are not there yet.

FP&A scenario modelling. AI-assisted scenario planning, sensitivity analysis, and rolling forecasts. The productivity benefit is real for FP&A teams spending significant time rebuilding models each planning cycle. The prerequisite is a financial model architecture that is consistent and documented well enough for an AI system to work with. Ad hoc models built in Excel by different analysts over different years are not a workable foundation.

THE EMERGING USE CASE WORTH WATCHING

Gartner identifies code generation as the highest-impact emerging use case in finance AI. This means using large language models to write Python scripts, SQL queries, or automation code that finance professionals can use without engineering support. For finance functions trying to build custom automation without engineering resource, this is worth taking seriously.

Agentic AI, where AI systems take autonomous actions across multiple tools, is at the opposite end of the maturity curve. Only 6% of finance leaders are currently using it, though 38% plan adoption within 12 months. Deploying agentic AI in any regulated finance process before the EU AI Act compliance obligations are clear is premature.

THE NUMBER THAT CUTS THROUGH THE NOISE

BCG's research puts median ROI from AI in finance at around 10%. Not 300%. Not transformational. 10%.

10%
median ROI from AI in finance (BCG research)

6%
of organisations qualify as AI high performers (McKinsey)

~85%
of AI projects fail to deliver intended value (Gartner, 2024)

That is not a reason to avoid AI investment. A 10% efficiency gain in a finance function processing thousands of transactions per month is meaningful. But it is a useful calibration against vendor presentations promising step-change returns. The organisations achieving above-median results are, without exception, the ones that did the foundational work first.

Which raises the obvious question: if the use cases are well understood and the potential returns are measurable, why do so many AI projects in finance fail to deliver? That is the subject of the next chapter.

CHAPTER 3: WHY MOST AI PROJECTS FAIL, AND IT IS NOT THE TECHNOLOGY

3

Between 70 and 85% of AI projects fail to meet their objectives. Only 48% make it into production at all. 30% of generative AI pilots are abandoned after the proof-of-concept stage, most commonly for one of four reasons: poor data quality, inadequate risk controls, escalating costs, or unclear business value.

These numbers are not about AI being immature technology. They are about organisations deploying capable technology onto inadequate foundations.

The technology, in most cases, works as advertised. The problem is what the technology is being asked to work with.

WHERE THE TIME ACTUALLY GOES

When an organisation undertakes an AI implementation in finance, the expectation is that the majority of the effort goes into configuring the AI system: training models, testing outputs, building interfaces. The reality is different. Data engineering, the work of finding, cleaning, standardising, and connecting data from across the organisation, consumes 60 to 80% of the total project timeline.

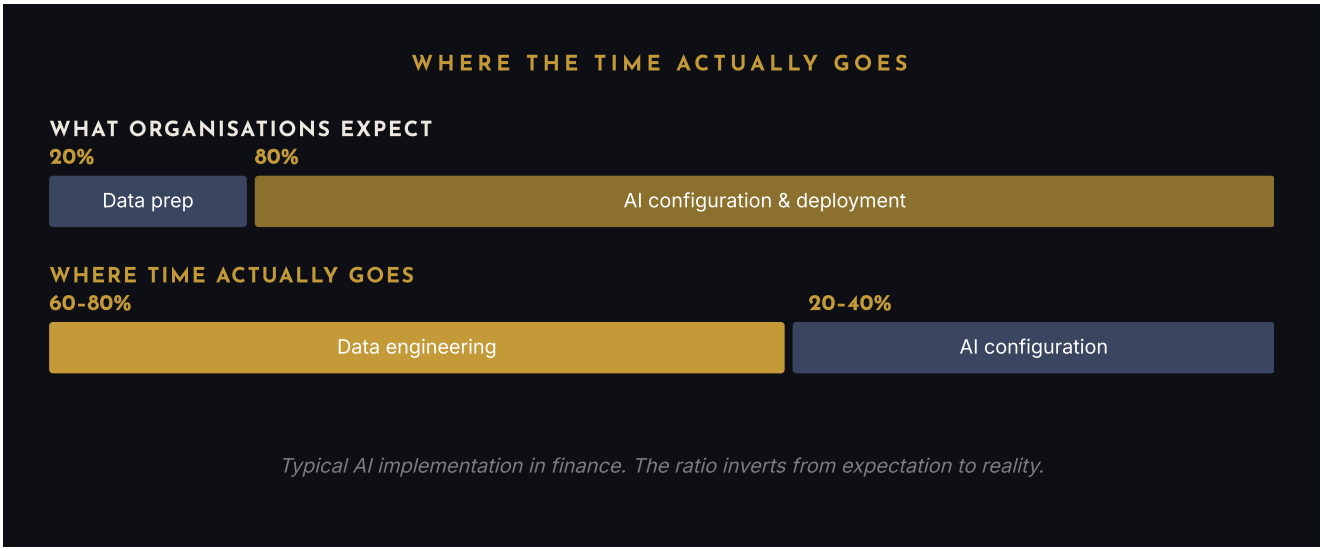

Typical AI implementation in finance. The ratio inverts from expectation to reality.

That figure surprises people who have not been through it. It should not. Finance data in most mid-market organisations is held across multiple systems that were not designed to work together: an ERP that predates the current management team, a CRM with its own revenue definitions, spreadsheets maintained by individuals who may no longer be with the business, and a chart of accounts that has been modified incrementally over years without a coherent structure.

Before any AI system can process this data, someone has to make it consistent. That work is slow, detailed, and not glamorous. It is also non-negotiable.

THE ROOT CAUSES, ONE BY ONE

01 · FRAGMENTED DATA

Multiple systems, each with its own data structure and coding conventions. Integration requires understanding both systems and finance processes.

02 · INCONSISTENT DEFINITIONS

Revenue means six different things depending on which report you read. AI cannot resolve definitional inconsistency - it reflects it back at scale.

03 · UNSTANDARDISED PROCESSES

You cannot automate a process that is different every month. Fixing it means documenting what is actually happening, then standardising.

04 · LEGACY SYSTEM CONSTRAINTS

AI tools sit on top of ERP systems, not instead of them. If the ERP data structure does not support the required outputs, the AI system fails.

05 · CULTURAL RESISTANCE

Finance teams that have been through failed technology implementations are not wrong to be sceptical. Change management is not a soft consideration: it is a project risk with a material probability of realisation.

THE INSIGHT THAT CHANGES THE FRAMING

THE CENTRAL INSIGHT

> Good accounting practice is AI readiness.

This is not a metaphor. Clean reconciliations, a consistent chart of accounts, documented month-end procedures, standardised journal entry templates, clearly defined intercompany policies: these are the foundations that determine whether an AI deployment succeeds or fails. They are also the foundations that determine whether your management accounts are reliable, whether your audit goes smoothly, and whether your board can trust the numbers they are seeing.

Poor data quality is estimated to cost organisations 15% of revenue annually. That cost exists independent of any AI project. Fixing it is not AI preparation: it is good finance practice that happens to also make AI viable.

The organisations treating AI readiness as a separate workstream, distinct from their core finance improvement agenda, are creating unnecessary complexity. The work is the same work. The sequencing is what matters.

WHY THIS WORK REQUIRES SPECIALIST INPUT

The foundational work described above is not work that a finance team can absorb alongside their existing responsibilities. Month-end close does not pause for a data quality project. Board packs still need to go out. VAT returns do not wait.

Assessing what needs fixing, sequencing the remediation, managing the change across the team, and building the governance framework that will govern AI use once it is deployed: this is transformation work, and it requires someone who has done it before.

> WORK WITH MAEBH
>
> The foundational work described in this chapter is the same work I have been delivering across finance teams for twenty years. Data quality, process documentation, system adoption, controls: these are not AI preparation tasks. They are finance function fundamentals that happen to be prerequisites for AI that works.
>
> If you want to talk through where your finance function sits against these foundations, that conversation is free.
>
> **MAEBHCOLLINS.COM/WORK-WITH-ME**

Chapter 4 sets out the AI-Ready Finance Framework: a five-pillar structure for assessing your current position and building the foundations that AI deployment actually requires.

CHAPTER 4: THE AI-READY FINANCE FRAMEWORK: FIVE PILLARS OF READINESS

Every framework in this space has a provenance. This one came from doing the work.

The five pillars below were not assembled from research reports or designed for a whitepaper. They emerged from the pattern of what consistently needed fixing before automation, systems integration, or AI deployment could succeed. When I eliminated hundreds of hours of monthly finance processing across multiple finance teams, the same categories of foundational work appeared every time, in every organisation, regardless of size or sector. Data. Process. Technology. People. Governance. The sequence varied. The categories did not.

Each pillar maps directly to a section of the AI Readiness Self-Assessment in Chapter 7. Use this chapter to understand what each pillar means. Use the scorecard to measure where your function currently stands.

PILLAR 1: DATA READINESS

Data readiness covers the quality, availability, consistency, and integration of the data your finance function produces and relies on. It asks whether your data is accurate enough to trust, structured consistently enough to use across systems, and accessible to the tools that need it.

> **DIAGNOSTIC QUESTION**
>
> If you pulled the same revenue figure from your ERP, your CRM, and your management accounts for the same period, would all three match?

NOT READY

Data exists in multiple systems with no single source of truth. Reconciliation between systems is a manual monthly task. Key metrics are defined differently depending on who produced the report.

READY TO PILOT

Core financial data is consistent across systems, reconciliation is automated or near-automated, and definitions are documented and applied uniformly.

PILLAR 2: PROCESS MATURITY

Process maturity covers whether your finance function's core processes are standardised, documented, and stable enough to build on. A process that is different every month depending on who is running it cannot be automated. A process that exists only in someone's memory cannot be handed over, improved, or audited.

DIAGNOSTIC QUESTION

If your most experienced finance team member left tomorrow, could the month-end close still run on time?

NOT READY

Key processes are undocumented, or documentation exists but does not reflect what actually happens. Close timelines extend when specific people are absent. Manual workarounds are embedded in the process rather than flagged as exceptions.

READY TO PILOT

Core processes are documented to a level that allows someone new to follow them. The close runs consistently regardless of who is in the team.

PILLAR 3: TECHNOLOGY FOUNDATIONS

Technology foundations cover the ERP infrastructure, cloud readiness, and API capability that AI tools connect to, draw data from, and write outputs back into. If the ERP is old, poorly configured, or unable to expose data through an API, the integration will fail or require expensive custom work.

DIAGNOSTIC QUESTION

Do you know which version of your ERP you are running, and when it was last meaningfully updated?

NOT READY

ERP is on a legacy version with significant customisations that complicate upgrades. Data extraction requires manual exports. No API integration capability exists or has been tested.

READY TO PILOT

ERP is current or recently updated, cloud-hosted or cloud-capable, and can expose data to external tools without manual intervention.

PILLAR 4: PEOPLE AND SKILLS

People readiness covers whether your finance team has the data literacy, analytical capability, and change readiness to work effectively with AI tools. The technology can be perfect and still fail at this layer. Finance teams that have been through failed technology implementations carry scepticism that is rational and earned. Addressing it requires more than a training session.

DIAGNOSTIC QUESTION

Could you name the person in your finance team who would champion an AI pilot, and would their colleagues trust them to do it?

NOT READY

The team has low confidence with data outside familiar reports. Previous technology changes were imposed rather than co-developed. There is no internal sponsor for change with credibility among the team.

READY TO PILOT

At least one person in the finance team has genuine analytical curiosity and the confidence to experiment. The team understands why change is happening, not just what is changing.

PILLAR 5: GOVERNANCE AND CONTROLS

Governance readiness covers the strength of your internal controls framework, the quality of your audit trail capability, and your awareness of the regulatory obligations that apply to AI use in finance. AI does not reduce the need for controls: it changes where the control points are. An AI system that processes journal entries needs the same segregation of duties, approval workflows, and exception reporting as a manual process.

DIAGNOSTIC QUESTION

If an AI system produced an incorrect journal entry that was posted and paid, how quickly would your current controls catch it, and who would be accountable?

NOT READY

Controls are informal or inconsistently applied. Audit trails are incomplete. There is no policy governing AI use in finance processes, and no awareness of the EU AI Act obligations that apply from August 2026.

READY TO PILOT

ICFR framework is documented and tested. Audit trail capability is robust. At least one person understands the regulatory obligations and has begun mapping AI use cases against them.

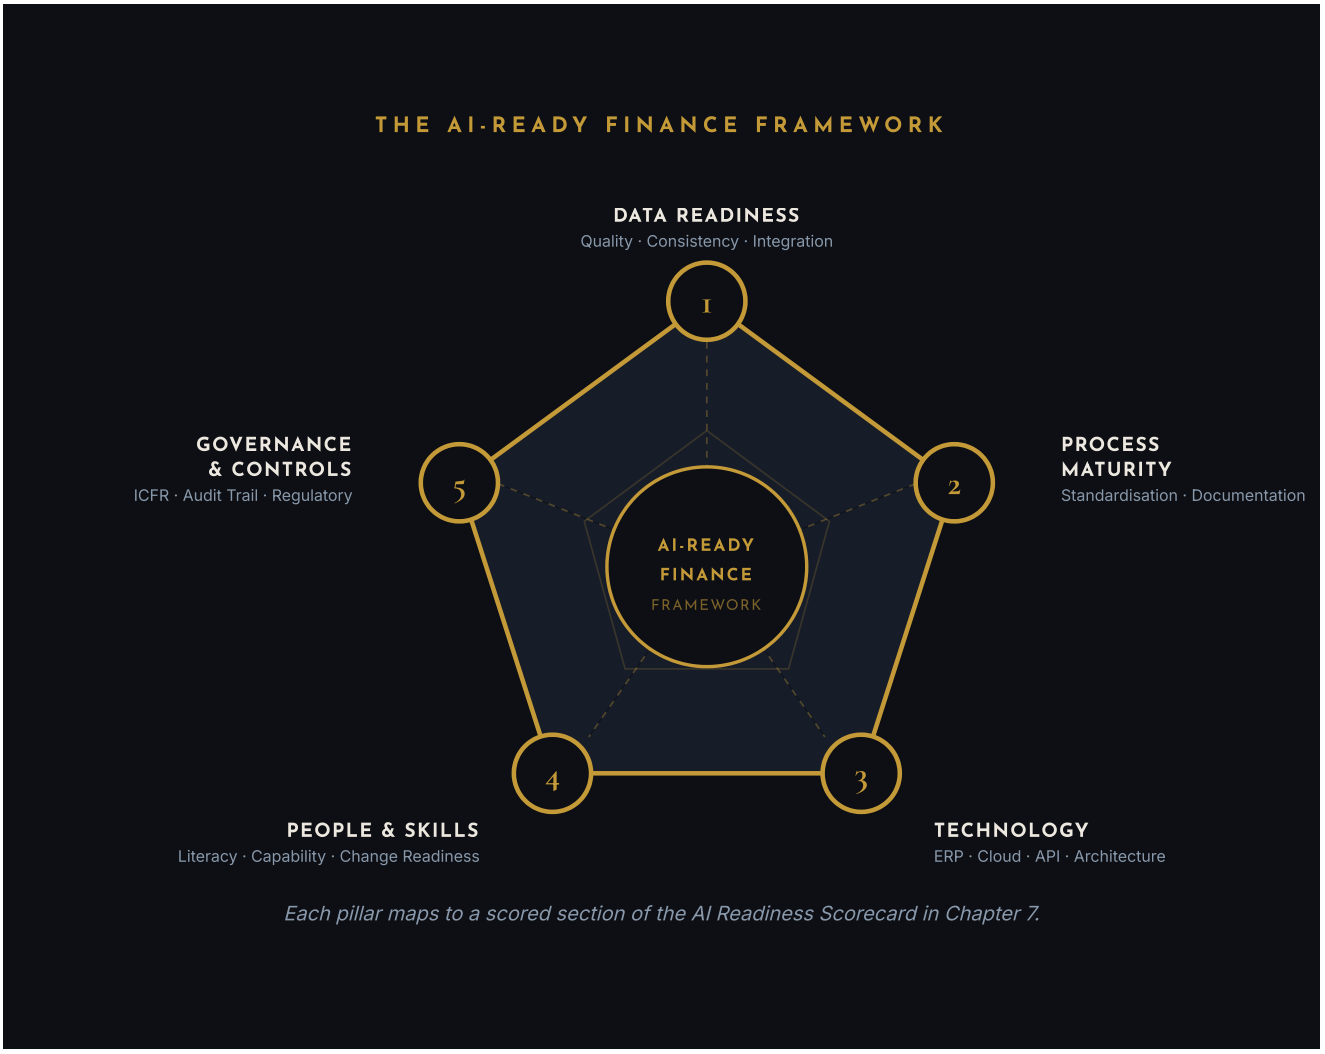

USING THE FRAMEWORK

No finance function scores at the top of all five pillars before embarking on AI deployment. The point is not perfection: it is informed sequencing. Knowing which pillars are weakest tells you where to focus first, which use cases are viable now, and where the deployment risk is highest.

Governance and controls, the fifth pillar, deserves particular attention in 2026. The regulatory obligations now applying to finance functions under the EU AI Act and the FRC's guidance on AI in audit are specific, consequential, and not yet widely understood. That is the subject of the next chapter.

CHAPTER 5: GOVERNANCE, REGULATION, AND AI: WHAT FINANCE LEADERS NEED TO KNOW IN 2026

5

The regulatory conversation about AI in finance has moved from future-tense to present-tense faster than most finance teams expected.

The EU AI Act (Regulation 2024/1689) entered into force in August 2024. Key obligations for high-risk AI systems and general-purpose AI in professional contexts take effect from August 2026. The deadline is August 2026. Many finance functions have not yet begun the mapping exercise required.

For most mid-market businesses, the immediate task is knowing where you stand, not building a dedicated compliance function. But knowing where you stand requires doing something, and very few finance teams have started.

WHAT THE EU AI ACT REQUIRES FROM FINANCE

The AI Act classifies AI systems by risk level. Finance use cases that attract the closest scrutiny include automated credit assessment, AI-assisted financial reporting, and AI tools used in audit preparation. For most mid-market finance functions, the practical obligations centre on three things: transparency (being able to explain what an AI system does and why it produced a given output), human oversight (a documented review process for AI-generated outputs before they become decisions or entries), and data governance (ensuring that the data the AI system uses meets defined quality requirements).

None of these are new concepts. They are the same principles that underpin a well-run ICFR framework. The difference with AI is that the obligation to document them is now statutory in many contexts, not just best practice.

AI RELOCATES CONTROLS. IT DOES NOT REMOVE THEM.

A manual reconciliation process has a control: the reviewer who checks the output. An AI-assisted reconciliation process has a different control: the human who reviews the AI output before it is posted, the exception report that flags what the AI could not match, and the audit trail that shows what the AI did

and when.

If your current ICFR framework was built before 2024, it almost certainly does not cover these control points. Most frameworks do not. This is a gap that needs to be identified and closed before deployment, not after.

THE AUDIT TRAIL PROBLEM

Auditors are asking questions about AI-generated outputs. The FRC's guidance on AI in audit, published in 2024, makes clear that audit firms are expected to understand the AI tools their clients use in financial reporting and to assess whether those tools introduce material risk.

In practice: if you are using AI to assist with journal entry processing, account reconciliation, or management account preparation, your auditor will want to know how the tool works, what oversight exists, and whether outputs can be independently verified. A finance function that cannot answer those questions clearly has an audit problem, not just an AI problem.

WHO IS ACCOUNTABLE WHEN AI GETS IT WRONG?

This question must be answered before deployment, not after. When an AI system produces an incorrect journal entry that is posted and paid, who is accountable?

In regulatory terms, the answer is clear: the organisation deploying the AI system owns the outputs. Vendor contracts routinely disclaim liability for AI errors. The governance framework must fill that gap.

Build a human review checkpoint for every AI-generated output that affects the financial statements. In practice, that means: a named reviewer who checks the AI output against the underlying data before posting; a sign-off field in the reconciliation or journal entry workflow that captures who reviewed it and when; and a documented escalation path for items the AI flagged as exceptions. The checkpoint does not need to be elaborate. It needs to exist, to be applied consistently, and to leave a trail. This is not slowing down AI adoption. It is making it defensible when something goes wrong, and something will go wrong.

HUMAN REVIEW CHECKPOINT

- ☐ **Named reviewer.** A specific person is accountable for checking the AI output against the underlying data before posting. Not a team. A named individual.
- ☐ **Sign-off field.** The reconciliation or journal entry workflow captures who reviewed it and when. This is the audit trail. It must exist within the output, not in a separate document.
- ☐ **Documented escalation path.** For outputs the AI flagged as exceptions, there is a defined process for what happens next. Not an informal understanding: a documented procedure.

GOVERNANCE IS WHAT MAKES THE OTHER FOUR PILLARS HOLD

Clean data, documented processes, a current ERP, and a capable team are necessary. They are not sufficient. A governance framework that does not account for AI-assisted processes is the gap that turns a successful pilot into an audit finding.

The self-assessment scorecard in Chapter 7 includes a dedicated governance section. The question to sit with before you complete it: if your external auditor asked to see your AI governance framework today, what would you show them?

> GOVERNANCE AND COMPLIANCE READINESS
>
> Most mid-market finance teams have not begun the governance mapping described in this chapter. The August 2026 deadline for EU AI Act obligations is closer than it looks, particularly if your organisation is at an early stage of the assessment.
>
> If you want a structured review of where your ICFR framework and AI governance documentation stand today, I can help with that work directly.
>
> **MAEBHCOLLINS.COM/WORK-WITH-ME**

Chapter 6 translates all five pillars into a 90-day action sequence.

6

CHAPTER 6: YOUR 90-DAY ACTION PLAN

Most AI readiness programmes fail for the same reason most finance transformations fail: too much designed at once, not enough delivered at all.

The 90-day plan below is not a transformation programme. It is a sequenced preparation: enough to move from uncertain to ready-to-pilot, across all five pillars, without derailing the finance function in the process. I have used this sequence in practice.

THE 90-DAY READINESS SEQUENCE

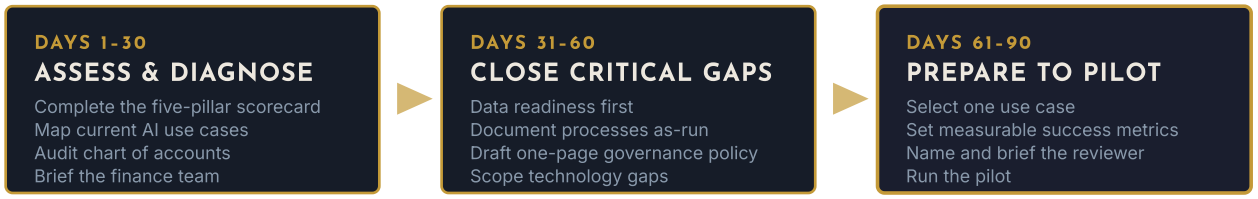

DAYS 1 TO 30: ASSESS AND DIAGNOSE

Before fixing anything, understand what you are working with. The self-assessment scorecard in Chapter 7 is the starting point. Complete it honestly, with your finance team, not in isolation. The gaps that emerge will shape everything that follows.

The specific tasks for this phase:

Complete the five-pillar scorecard. Score each pillar, note the evidence for each score, and identify the two or three lowest-scoring areas. These are your critical path items.

Map your AI use cases. List every AI tool your finance team is currently using or piloting, including general-purpose tools like ChatGPT or Copilot that may not be formally sanctioned. Assess each against the governance pillar criteria from Chapter 5.

Audit your chart of accounts. Check for inconsistency across entities, redundant codes, and categories that cannot be reliably mapped to a standard reporting structure. This is the single most common data readiness failure and the one that blocks the most AI use cases.

Brief your team. Tell them why you are doing this assessment. Finance teams that are informed and involved in readiness work are significantly more likely to adopt AI tools effectively than those for whom it arrives as a done deal.

By the end of day 30, you should have a completed scorecard, a map of current AI usage, a list of the three to five highest-priority gaps, and a team that understands what is coming.

DAYS 31 TO 60: CLOSE THE CRITICAL GAPS

Do not attempt all five pillars simultaneously. Address the gaps that would prevent any AI pilot from succeeding, in this order:

Data readiness first. If the chart of accounts is inconsistent, fix it. If reconciliation data lives in multiple systems with no integration, document the manual bridge while you plan the technical fix. You do not need perfect data to begin piloting: you need data that is consistent enough to produce a reliable baseline.

Process documentation second. Identify the two or three finance processes most likely to be targeted for AI assistance. Document them as they actually run, not as they are supposed to run. The gap between these two is usually where the most risk lives.

Governance framework third. Draft a one-page AI governance policy covering: what AI tools are approved for use in finance processes, what human review is required before AI outputs are acted on, and who is accountable for AI-assisted outputs.

It does not need to be comprehensive. It needs to exist.

If technology foundations are a significant gap (legacy ERP, no API capability), this phase is also the time to scope that work separately. A full ERP upgrade will not happen in 60 days, but understanding the timeline and the constraints will allow you to select appropriate AI use cases.

DAYS 61 TO 90: PREPARE TO PILOT

By day 60, you have a clearer picture of your readiness and your constraints. Days 61 to 90 are about selecting one use case and preparing to run it properly.

Select the pilot use case based on two criteria: the pillar scores suggest it is viable, and the outcome is measurable. Reconciliation automation, intelligent document processing for accounts payable, and management account variance commentary are all well-suited first pilots for mid-market finance functions.

Set the success metrics before the pilot starts. Define what good looks like: time saved, error rate, reviewer confidence.

Without a pre-agreed baseline, the pilot will produce opinions rather than evidence.

Identify the reviewer. Name the person responsible for checking AI outputs during the pilot. Confirm they have the authority to reject an AI output and the time to actually review it.

The 90-day plan does not end with an AI deployment. It ends with a finance function that knows what it is ready for, has closed its most significant gaps, and has a governance framework in place to run a pilot safely. That is a materially better position than most mid-market finance functions are currently in.

The self-assessment scorecard in Chapter 7 is designed to be completed at the start of this process and revisited at day 90. The change in scores is the measure of progress.

CHAPTER 7: AI READINESS SELF-ASSESSMENT SCORECARD

7

This scorecard maps directly to the five pillars introduced in Chapter 4. Use it at the start of your 90-day plan to establish a baseline, and again at day 90 to measure progress. Complete it with your finance team, not in isolation: the scores are only useful if they reflect what is actually true, not what you would like to be true.

Score each statement from 1 to 4:

1 = Not in place **2** = Partially in place **3** = Largely in place **4** = Fully in place

PILLAR 1: DATA READINESS

STATEMENT	1	2	3	4
Our chart of accounts is consistent across all entities and has not been significantly modified in the last 12 months				
Core financial data (general ledger, AP, AR, payroll) sits in a single system or is reliably reconciled between systems				
Key metrics are defined consistently and produce the same result regardless of who pulls the report				
Data quality issues are identified, logged, and owned by a named person				
PILLAR 1 TOTAL: /16				

PILLAR 2: PROCESS MATURITY

STATEMENT	1	2	3	4
Core finance processes (month-end close, reconciliations, AP, payroll) are documented to a level that allows someone new to follow them				
The month-end close runs to the same timeline regardless of who is in the team				
Manual workarounds are identified and logged as exceptions, not embedded as standard practice				
Process documentation is reviewed and updated at least annually				
PILLAR 2 TOTAL: /16				

PILLAR 3: TECHNOLOGY FOUNDATIONS

STATEMENT	1	2	3	4
Our ERP is on a current or recently updated version				
Financial data can be extracted to external tools without manual intervention				
We have tested or use at least one API integration between finance systems				
Someone in the finance team can specify what data our ERP can expose via API and what it cannot				
PILLAR 3 TOTAL: /16				

PILLAR 4: PEOPLE AND SKILLS

STATEMENT	1	2	3	4
At least one person in the finance team can work confidently with data outside of standard reports				
The team understands why AI is being considered and has been consulted on it				
The finance team is actively involved in decisions about technology change, not informed after the fact				
There is a named internal sponsor for finance technology change with credibility among the team				
PILLAR 4 TOTAL: /16				

PILLAR 5: GOVERNANCE AND CONTROLS

STATEMENT	1	2	3	4
Our ICFR framework is documented, tested, and current				
Audit trails for financial transactions are complete and independently verifiable				
We have a policy (however brief) governing the use of AI tools in finance processes				
We have mapped our current and planned AI use cases against EU AI Act obligations				
PILLAR 5 TOTAL: /16				

YOUR TOTAL SCORE: /80

SCORE	READINESS LEVEL	RECOMMENDED ACTION
20 or below	Significant gaps	Begin with diagnosis and foundational work. Do not evaluate AI vendors until the data and process pillars are addressed.
21 to 40	Partial readiness	Specific pillars need attention before piloting. Use the 90-day plan in Chapter 6. Low-risk use cases may be viable with robust governance.
41 to 60	Moderate readiness	Ready to pilot in well-defined, lower-risk use cases. Governance and human oversight are essential at this stage.
61 to 80	Strong readiness	Proceed to pilot selection. Focus on change management and scaling beyond the first use case.

YOUR AI READINESS PROFILE

Score each pillar using the scorecard above, then mark your score on each axis and connect the dots. Each ring represents one point (1 to 4).

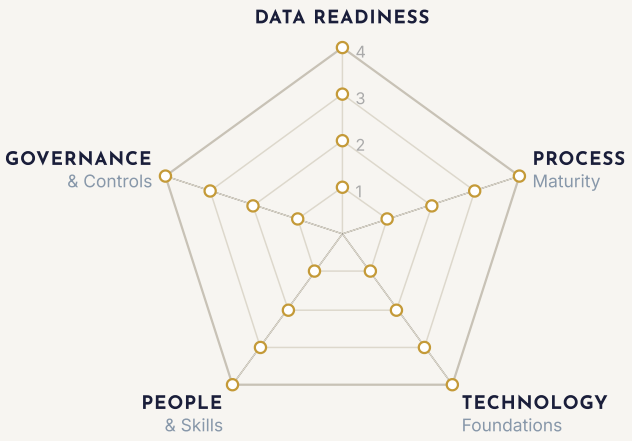

1 = Not in place · 2 = Partially · 3 = Largely · 4 = Fully in place

HOW TO USE THIS SCORECARD

A low overall score is not a reason to delay indefinitely. It is a map of where to start. Address the lowest-scoring pillar first, not the most interesting one.

Rescore at day 90. The change in scores is more informative than the absolute numbers. A finance function that moves from 28 to 52 in 90 days has done meaningful preparatory work. One that scores 65 at the start but has not addressed governance is less ready than the score suggests.

The scorecard is a tool, not a verdict. Use it to focus effort, not to justify inaction.

READY TO ACT ON YOUR SCORES?

If you have just scored your finance function and identified specific gaps, the next step is clear: address the lowest pillar first, not the most interesting one.

If you want that work done properly, with someone who has delivered it before, a readiness assessment is where most engagements start. It takes two to four weeks, works against your actual data and processes rather than self-reported scores, and produces a prioritised plan that tells you exactly where to focus and in what order.

MAEBHCOLLINS.COM/WORK-WITH-ME · HELLO@MAEBHCOLLINS.COM

ABOUT THE AUTHOR

Maebh Collins is a Fellow Chartered Accountant (FCA, ICAEW) with Big 4 training and twenty years of finance leadership across commercial businesses, international operations, and public sector environments.

She has founded and led two international businesses, delivered finance transformation across multiple organisations, and implemented the ERP systems, automation frameworks, and controls structures that made those transformations hold. The technical and the operational are not separate disciplines in her work. When off-the-shelf solutions did not exist, she built the integrations herself.

Her background in Big 4 corporate tax and transfer pricing, including a £125m tax burden reduction delivered for a major UK PLC, gave her a rigour for evidence and precision that she has applied to every finance function she has led or transformed since. The difference between advising on difficult decisions and making them with your own money at stake has shaped how she approaches every engagement.

She writes on AI in finance, finance transformation, and CFO leadership. Her work is published at maebhcollins.com.

WHAT COMES NEXT

You have your scores. You know which pillars need work. That is the useful part.

The question most finance leaders get wrong at this point is sequencing. They try to address every pillar at once, treat the lowest score as a verdict rather than a starting point, and begin evaluating vendors before the foundational work is done. A lower score addressed systematically produces better outcomes than a higher score left unexamined.

AI READINESS ASSESSMENT

A structured diagnostic across all five pillars, conducted against your actual data, processes, and systems: not self-reported scores alone. Includes a prioritised remediation plan and a clear answer to whether you are ready to pilot, and if not, what needs to happen first.

Typical engagement: two to four weeks.

FINANCE TRANSFORMATION

Where the diagnostic identifies foundational gaps: process documentation, data quality, ERP configuration, or controls frameworks, I work alongside your finance team to close them. Every piece of this work pays back independently of any AI deployment.

Engagements from three months.

ADVISORY AND INTERIM LEADERSHIP

For businesses that need senior finance capability on a sustained basis: as a CFO, Finance Director, or advisory role throughout a transformation programme.

If you scored your finance function and want a second opinion on what the numbers mean, or if you want to discuss where to start, get in touch. That conversation costs nothing.

Work With Me | **hello@maebhcollins.com**

GLOSSARY

Agentic AI: AI systems that can take sequences of actions autonomously, such as retrieving data, running calculations, and drafting outputs, without step-by-step human instruction. Relevant to finance in the context of automated reporting and workflow execution.

API (Application Programming Interface): A defined connection point that allows two software systems to exchange data. Whether your ERP has API capability determines whether AI tools can draw on live financial data or require manual data exports.

EU AI Act (Regulation 2024/1689): The European Union's regulatory framework for artificial intelligence, in force from August 2024. Key obligations for high-risk and general-purpose AI systems take effect from August 2026.

Hallucination: The tendency of large language models to generate plausible-sounding but factually incorrect outputs. An LLM asked to summarise a management account may produce a figure that does not exist in the underlying data, presented with the same confidence as one that does. In financial reporting, an undetected hallucination is not a technical curiosity: it is a misstatement. Human review of AI-generated financial outputs is not optional for this reason.

ICFR (Internal Controls over Financial Reporting): The framework of policies, procedures, and oversight mechanisms that govern the accuracy and reliability of financial reporting. The primary control structure that AI governance policies must integrate with.

IDP (Intelligent Document Processing): AI-powered extraction, classification, and processing of information from documents such as invoices, contracts, and bank statements. One of the most practical near-term AI applications for mid-market finance functions.

LLM (Large Language Model): The class of AI model that underlies tools such as ChatGPT, Copilot, and Claude. Trained on large volumes of text to generate, summarise, and analyse language. Not inherently a finance tool: its value in finance depends entirely on the quality of the data and processes it is applied to.

RAG (Retrieval-Augmented Generation): A technique that connects an LLM to a specific body of documents or data, allowing it to draw on that source rather than relying solely on its training data. Relevant for finance use cases involving internal policies, procedures, or proprietary datasets.

AI-Ready Finance

The Practitioner's Guide for UK and European Mid-Market Finance Teams

WORK WITH MAEBH

maebhcollins.com/work-with-me

hello@maebhcollins.com

Maebh Collins FCA

Fellow Chartered Accountant · Big 4 Trained · Finance Transformation Specialist