

Prospero

Order out of complexity.

The Intelligent Firm

Why AI Will Redraw the Operating Model
for Fund Administrators and Corporate
Service Providers – and What Leadership
Teams Should Do Now

Independent advisory and operational
leadership for the offshore financial
services market.

 prosperoadvisory.je

 paul@prosperoadvisory.je

 Jersey, Channel Islands

EXECUTIVE BRIEFING

The Intelligent Firm

Why AI Will Redraw the Operating Model for Fund Administrators and Corporate Service Providers – and What Leadership Teams Should Do Now

A six-minute read for boards, CEOs and COOs.

The full analysis, evidence and Jersey regulatory detail are in the deep-dive annex that follows.

Prepared by

Prospero Advisory Limited

Governance & Operational Advisory | Jersey, Channel Islands | July 2026

Declared up front: Prospero also builds AI tools for this sector (Prospero Tools). This briefing draws on both advisory and practitioner experience; weigh it with that interest in view.

Executive summary

The operating model that built the trust, corporate services and fund administration industry, growth through headcount, is reaching its limit. Margins have been compressed by a decade of regulatory expansion, and capacity is still bought one hire at a time. **Artificial intelligence breaks that link between revenue and headcount**, and it does so with tools that are deployable today.

The opportunity concentrates in five areas: onboarding and KYC/AML, entity administration, fund administration, regulatory reporting and client communication. Across all five the pattern is identical, AI absorbs the procedural volume, professionals keep judgement, accountability and the client relationship.

This is not speculative. **81% of financial services firms are already adopting AI**, 52% are actively adopting *agentic* AI that orchestrates entire workflows. Yet only 6% of traditional institutions have reached the transforming stage against 19% of fintechs, and fintech-native administrators are entering the market with structurally lower cost bases. For a Jersey firm the regulatory wind is favourable: the JFSC issued dedicated guidance on the use of AI in July 2026 (five principles, no new requirements, and explicit encouragement to explore the technology), uses AI in its own operations, and a 150% RegTech tax super-deduction improves the investment case.


The numbers are modelled, not asserted. **Our worked example, a 40-person operations function costed on published Channel Islands salary data, releases the equivalent of nine professionals in the base case: approximately £417,000 net per year, recurring.** The full model, assumptions and sensitivities are in the annex.

The strategic choice: **Defend the margin** (deliver the same output with a leaner team) or **fund the growth** (scale without hiring in proportion), or a deliberate blend of the two, decided on evidence. The diagram below shows both. Either way, the window for first-mover advantage is open now and will not stay open indefinitely. One caution belongs alongside it, hours compressed under time-based billing reduce revenue before they reduce cost. The pricing model must move with the operating model, and the annex addresses both together.

THE ASK *A focused pilot in two or three high-value processes, wrapped in proper governance, will prove the model within a single financial year. The firms that move with purpose now will build cost advantages their competitors cannot easily replicate.*

The AI-augmented operating model

The model that emerges is not, in the near term, one with dramatically fewer people. It is one where the professional workforce does different and more valuable work. An AI capability layer sits beneath the professionals, absorbing high-volume procedural work, with human review retained over every regulated output. That single structural change supports two distinct strategies:



The AI capability embeds as a new operating layer beneath professional judgement. Left: compress the cost base. Right: grow without hiring in proportion.

Panel A: Defend the margin. A function that today needs 50 people may need around 38 in our modelled base case, and around 35 in the stretch case shown. For a firm generating £30m of revenue, a sustained efficiency gain of that order is transformative.

Panel B: Fund the growth. The same released capacity can instead be redeployed into new business. Revenue rises while headcount stays broadly flat, the operating leverage that headcount-bound competitors cannot match.

Crucially, the regulatory framework in jurisdictions like Jersey requires humans to remain in the loop for regulated activity. That is not a constraint on the model, it is a feature of it, and one to communicate proactively to regulators and clients alike.

Where the opportunity concentrates

The cost base of these businesses is dominated by procedural work (document review, drafting, data extraction, routine correspondence), not by advisory judgement. That is precisely the work AI transforms first and most completely. The map below links each area to the mechanism that changes it.

WHERE THE TIME GOES — AND WHERE AI CHANGES IT

The five process areas that dominate cost in fund administration and corporate services

PROCESS AREA	DOMINANT WORK TODAY	AI MECHANISM	EFFECT
Client onboarding & KYC/AML	Document review, data extraction, UBO & screening, CDD file build	Document intelligence + contextual screening + draft risk scoring	Days → hours
Entity admin & corporate governance	Drafting resolutions, minutes, registers, deadline tracking	LLM drafting in house style + portfolio deadline monitoring	Draft in seconds
Fund administration	Investor queries, side-letter tracking, FATCA/CRS, reporting	Automated query handling + obligation extraction & checks	Most queries auto
Regulatory reporting	Recurring returns following predictable, rule-based patterns	Templated preparation from well-structured source data	Substantially auto
Client communication	Status updates, checklists, routing, internal coordination	Drafted responses + status automation with human sign-off	Overhead removed


The common pattern: the majority of time spent is procedural, not advisory — and it is the procedural work that AI transforms first.

Each of the five cost-dominant process areas maps to a specific AI mechanism and a measurable operational effect.

The frontier: the AI-native firm

The gains above are available to any well-governed firm within its existing architecture. A further level exists for firms prepared to rebuild around the technology. **The economics of software creation have collapsed:** prompt-native development platforms let the people who know the workflows best build production systems internally, within CTO/CISO guardrails, and set a roadmap to retire the third-party SaaS estate, contract by contract. Staff redeploy from processing into building and operating AI agents: new roles (agent builders, operators, quality auditors and their trainees) filled overwhelmingly from the existing workforce, with no coding

background required. Over years, the client-facing team becomes smaller and paradoxically more available to clients.



Levels 1 and 2 optimise the firm you have. Level 3 rebuilds the firm around the technology.

We state the prize as conviction, not research. Headcount-equivalent gains of 50% and beyond, whichever route was chosen at Level 2, are in our view not only possible but where this industry ultimately lands. No regulated administrator has yet published the evidence, likely because no one has yet walked the path. The firms that walk it first will own that evidence. The full argument, including what cannot be built and the governance framework that makes the rest safe, is in the annex.

What leadership should do now

The firms that benefit most are not those launching the most ambitious transformation programmes, but those that prove high-value use cases quickly, build internal confidence and expand from evidence. A sensible first year has three moves:

- **Select two or three pilots** where the work is high volume, the output is objectively reviewable and errors are recoverable. KYC document review, board minute drafting and regulatory deadline monitoring are strong candidates.
- **Wrap them in governance from day one.** An AI-use policy, a data-handling protocol, an AI register, defined human-review points backed by sampling-based QA, a model-availability fallback and an incident process. In a regulated context this is non-negotiable, and it maps directly onto the JFSC's guidance on the use of AI issued July 2026, and the FSB's newly consulted sound practices.
- **Measure rigorously, then expand** into adjacent processes and decide the operating model question (defend the margin, fund the growth, or a blend) on evidence rather than assumption.

The professionals in these businesses are too well trained, and too expensive, to spend their days on work a machine now does better. The intelligent firm is the one that acts on that, and it is achievable today.

Read the full analysis. This briefing is a summary. The deep-dive annex that follows sets out the process-by-process detail, the modelled worked example, the AI-native frontier case, the Jersey regulatory landscape, the risk framework, the adoption roadmap, and a fully sourced reference appendix with clickable links.

Paul Monahan | Founder | Prospero Advisory Limited | paul@prosperoadvisory.je | +44 7797 922 174

prosperoadvisory.je | prosperotools.com | Jersey, Channel Islands

This briefing is provided for information only and does not constitute legal or regulatory advice.

The Intelligent Firm

How Artificial Intelligence Will Transform the Operating Model for Fund Administrators and Corporate Service Providers

This annex follows the Executive Briefing of the same title. It contains the full analysis, evidence, Jersey regulatory detail and sources.

Prepared by

Prospero Advisory Limited

Governance & Operational Advisory | Jersey

July 2026

Declared up front: Prospero advises on the transformation this paper describes and builds AI tools for the same sector (Prospero Tools). The analysis draws on both advisory and practitioner experience; readers should weigh it with that interest in view.

Executive Summary

The trust, corporate services and fund administration industry is at a crossroads. Decades of growth built on headcount expansion, manual process, and institutional inertia have produced firms that are increasingly expensive to operate, slow to scale and vulnerable to competitive disruption. Artificial intelligence represents not an incremental improvement, but a fundamental reimagining of how these businesses work.

This paper argues that AI, deployed thoughtfully and with appropriate governance, will compress the cost base of fund administrators and corporate service providers (CSPs), accelerate delivery, reduce operational risk and free the professionals within these firms to focus on genuinely high-value advisory work. The transition will require investment, cultural change and careful regulatory navigation, and the firms that treat AI as a strategic priority, rather than an IT upgrade, will emerge with operating models that are materially more competitive than those of today.

The scale of adoption is no longer in question. The 2026 Global AI in Financial Services Report produced by the Cambridge Centre for Alternative Finance in partnership with the BIS, IMF and World Economic Forum, found that 81% of financial services firms are already adopting AI at some level¹, and a concurrent Finastra survey found that only 2% of financial institutions report no use of AI at all². The frontier has already moved beyond assistive tools: 52% of firms are actively adopting *agentic* AI (systems that orchestrate entire workflows rather than draft individual documents), with 23% already at the scaling or transforming stage.³

Five operational areas will see the most significant early transformation: client onboarding and KYC/AML; entity management and corporate governance; fund administration; regulatory reporting; and client communication. Across all five, the pattern is the same, AI handles volume, pattern, and procedure. Professionals handle judgement, accountability and client relationships.

The economics are material and measurable. A worked example developed for this paper (a 40-person operations function at a small to mid-size Jersey administrator, built on published Channel Islands salary data) indicates that a realistic base case releases approximately 23% of team capacity at full adoption: the equivalent of nine full-time professionals, worth approximately £480,000 a year gross, or £417,000 net of tooling and governance costs.⁴ A conservative case releases 15% and a stretch case consistent with the best documented

¹Cambridge Centre for Alternative Finance, in partnership with the Bank for International Settlements, IMF and World Economic Forum, 2026 Global AI in Financial Services Report. Link in the Appendix.

²Finastra, Financial Services State of the Nation Survey (1,509 senior executives across 11 markets). Link in the Appendix.

³2026 Global AI in Financial Services Report: see note 1.

⁴Salary midpoints from Alexander Daniels Offshore, 2025 Channel Islands Salary Guide (link in the Appendix). The full assumption set, scenario logic and sensitivity analysis are set out in the companion model: The Intelligent Firm AI Cost Impact Model (Excel), available on request.

deployments releases 30%. Beyond these modelled gains, Section 5 sets out the frontier case, the AI-native firm, stated explicitly as conviction rather than research.

The firms that move with purpose now will build structural cost advantages that are difficult to replicate. The firms that wait may find the window for building that advantage has already closed.

A firm that compresses its cost base by a fifth and redeploys that capacity into business development will generate revenue growth that widens the competitive gap further. Meanwhile, new entrant boutiques, unburdened by legacy systems are beginning to offer AI-native propositions at materially lower price points. Section 4.2 addresses the question this raises for every board: when AI compresses the hours, the pricing model, not just the operating model, must move.

The intelligent firm is not a distant aspiration. It is achievable today with tools already available in the market. What it requires is strategic clarity, governance discipline and the organisational will to challenge long-established ways of working.

1. The Industry at an Inflection Point

The fund administration, trust, and corporate services industry has always been a people business. The value delivered to clients (regulatory compliance, entity governance, investor administration, AML/KYC oversight) has historically depended on trained professionals doing complex, judgement-intensive work. That model is not going away but it is about to change profoundly.

The past decade has brought well-documented pressure. Regulatory expansion in the wake of the global financial crisis, the OECD's economic substance regime, FATCA/CRS, and the continuing industrialisation of AML/CFT/CPF obligations, have each added operational cost, predominantly headcount cost, without a proportionate increase in what clients are willing to pay. Margin compression has been the defining commercial theme for large providers and boutiques alike for a decade.

Large players (Apex Group, IQ-EQ, Vistra, Alter Domus, and their peers) have responded through acquisition and near-shoring, spreading cost across larger revenue bases and building operational hubs in lower-cost jurisdictions. This has been a successful strategy for those who can execute it, but it has not solved the underlying problem. The operating model remains fundamentally manual, and the relationship between revenue and cost remains uncomfortably linear. Headcount continues to be the primary lever of capacity.

Artificial intelligence changes this equation. Large language models and the broader ecosystem of AI tooling now make it possible to automate, accelerate and augment many of the processes that have historically consumed most of the billable hours in these businesses. This is not a theoretical prospect awaiting future technology. The mid-2026 generation of frontier models extended what was already possible: the leading systems can now sustain autonomous, multi-step work over long horizons (reviewing a document set, chasing

exceptions and assembling an output for human approval) rather than merely responding to individual prompts.⁵

Crucially for Jersey-based businesses, the regulatory wind is at their backs. The JFSC's 2026–2030 strategy explicitly commits to encouraging and enabling responsible innovation and technology adoption, including AI, across its regulated community.⁶ The Government of Jersey's Competitiveness Programme reinforces this direction, and a 150% tax super-deduction on RegTech expenditure, introduced in the 2024 Budget, provides material financial support for firms ready to invest.⁷ Sections 6 and 7 of this paper address the regulatory landscape and the management of AI risk in full.

The question facing boards and leadership teams is not whether AI will affect their business. It will. The question is whether their organisations will lead that change or respond to it.

2. The Current Operating Model: Where Time Actually Goes

To understand the scale of the AI opportunity, it is worth being precise about where time and cost are concentrated in the current model. Across most firms in this sector, the answer is depressingly predictable.

Client Onboarding and KYC/AML

New client onboarding in the fund administration and trust context is extraordinarily labour-intensive. For a typical fund client, the process involves collection and verification of constitutional documents, identification and verification of all underlying beneficial owners, directors, and associated persons, PEP and sanctions screening, assessment of source of funds and source of wealth, completion of an internal CDD file, and a risk assessment leading to a risk rating.

For a complex structure (a multi-class fund with an institutional investor base, or a complex corporate group with layered ownership), this can run to hundreds of hours of professional time. Much of that time is spent on document review, data extraction and form completion: precisely the tasks that AI handles well. Industry benchmarking puts the cost of a single corporate KYC review at approximately US\$2,500, before the opportunity cost of the professional time it consumes.⁸

Ongoing Monitoring

Regulators expect that KYC/AML compliance is not a one-time exercise but an ongoing programme. Periodic review cycles, trigger events, and continuous PEP and sanctions re-

⁵Anthropic released Claude Fable 5, the first publicly available Mythos-class frontier model, on 9 June 2026. The model was suspended for 19 days after security researchers demonstrated a jailbreak and redeployed on 1 July 2026 with strengthened safeguards. Links in the Appendix.

⁶Jersey Financial Services Commission, 2026–2030 Strategy, published March 2026. Link in the Appendix.

⁷Government of Jersey, Budget 2024: 150% super-deduction on qualifying regulatory-technology (RegTech) expenditure by financial services companies. Link in the Appendix.

⁸Fenergo, Financial Crime Industry Trends 2025 (use of advanced AI tools in KYC/AML rising from 42% of institutions in 2024 to 82% in 2025; periodic KYC review automation averaging roughly one-third) and Fenergo KYC cost benchmarking (average corporate KYC review approximately US\$2,500–2,600). Links in the Appendix.

screening create a rolling compliance burden that, at scale, requires significant dedicated resource. This is largely procedural work: high volume, rule-driven and repeatable.

Entity Administration

For a trust company or CSP managing a large book of entities, corporate governance administration is a constant activity. Drafting resolutions, maintaining statutory registers, coordinating director availability, preparing board packs and ensuring filing deadlines are met. Each individually manageable, but collectively a significant operational overhead at scale. The consequences of failures here, whether a missed filing or an unsigned resolution, can be material.

Fund Administration

The core fund administration workflow (NAV calculation, investor record-keeping, distribution processing, capital call and drawdown administration) has been partially automated for many years through dedicated platforms. However, the surrounding work is not. Investor queries, onboarding correspondence, FATCA/CRS classification, side letter tracking and bespoke reporting remain largely manual and disproportionately time-consuming relative to their actual complexity. It's hard to believe in 2026, but there are firms out there still using a mail merge to produce investor notices.

Regulatory Reporting

The compliance and reporting landscape has expanded dramatically. FATCA, CRS, DAC6, economic substance returns, AIFMD Annex IV reporting and the requirements of frameworks such as DORA create a recurring compliance calendar that is both technically demanding and heavily procedural. Much of the work involved in preparing and submitting these returns follows predictable patterns, precisely the conditions in which AI performs best.

Client Communication

A surprisingly large proportion of professional time is spent on routine client communication such as answering status queries, sending document requests, providing update emails and coordinating between internal teams. This communication overhead is largely invisible in billing metrics, and largely automatable.

KEY INSIGHT *Across these five areas, a common pattern emerges: the majority of the time spent is procedural, not advisory. It is the procedural work that AI will transform first and most completely.*

3. Where AI Changes Everything

3.1 KYC/AML and Client Onboarding

The transformation of KYC/AML through AI is no longer on the horizon, it is here today and some firms are already ahead. Fenengo's Financial Crime Industry Trends 2025 report found that the use of advanced AI tools in AML and KYC operations jumped from 42% of institutions in 2024 to 82% in 2025. Yet only around a third of periodic KYC review cycles are automated even at firms that have adopted AI. This is evidence that recognition has outpaced effective deployment, and that the majority of the efficiency opportunity in this area remains uncaptured.⁹ For fund administrators and CSPs, the gap between what is already technically achievable and what is actually embedded in daily operations is the most significant short-term opportunity.

Document intelligence tools can extract and classify information from constitutional documents, passports, corporate registry extracts and UBO declarations with high accuracy, eliminating the manual data entry that today consumes a significant share of compliance team time. AI-powered adverse media screening moves beyond keyword matching to contextual analysis, dramatically reducing false-positive rates while improving genuine hit detection. Risk scoring models, trained on portfolio data, can generate initial CDD assessments that professionals review and approve, rather than draft from scratch.

The practical effect: an onboarding that currently consumes days of professional effort, can now be compressed to hours. Elapsed time still depends on how quickly clients supply documents (AI won't fix client responsiveness!), but every document received is processed the day it arrives and nothing sits in a queue. For a firm onboarding 50 or 100 new structures every year, the cumulative saving is material, not as a one-off gain but as a permanent structural improvement to the cost of new business.

3.2 Entity Management and Corporate Governance

Large language models are particularly well suited to the drafting work that consumes significant time in the entity management context. Board minutes, resolutions, consent letters, powers of attorney, and routine correspondence follow established patterns, and AI can generate high-quality first drafts in seconds, leaving professionals to review, amend and sign off rather than write from scratch.

AI can also function as an intelligent governance management system: monitoring filing deadlines across a portfolio of entities, flagging upcoming requirements, generating draft filings and escalating exceptions for human review. The compliance assurance value, ensuring nothing falls through the cracks in a large entity portfolio, may be as commercially significant as the time saving itself.

3.3 Fund Administration

While core NAV processing and investor accounting are already technology-driven in most fund administration businesses, the surrounding workflow is not. AI will change several components of this.

⁹Fenengo: see note 8.

Investor query handling could be substantially automated. A well-designed AI system, with access to the fund's governing documents, financial statements and investor records, can answer most routine investor queries (balance enquiries, distribution history, tax document requests) without human involvement. Queries requiring genuine judgement can be escalated with a draft response already prepared, reviewed, and ready to send.

Side letter management, today a largely manual tracking exercise, becomes an automated compliance check. AI reads the side letter, extracts the obligations, tracks their fulfilment and alerts administrators when action is required. Regulatory reporting preparation (AIFMD Annex IV, FATCA/CRS, investor tax certificates) can similarly be substantially automated where the underlying data is well structured.

3.4 Compliance Monitoring and Regulatory Intelligence

AI-powered regulatory monitoring tools can track changes to legislation and regulatory guidance in relevant jurisdictions, assess their impact on a firm's client portfolio and generate briefing notes or action plans for professional review. For a Jersey-regulated business administering structures across multiple jurisdictions, this capability (essentially a continuously updated regulatory horizon-scanning function) has significant operational and commercial value. Its relevance is immediate, the JFSC's revised AML/CFT/CPF Handbook took effect on 31 May 2026, with further changes to MLCO requirements following from July, exactly the kind of rolling regulatory change this capability can be designed to absorb.

Transaction monitoring for AML purposes is another area of substantial AI capability. Pattern recognition across payment flows, counterparty networks and transaction characteristics are areas where AI materially outperforms rules-based systems, both in detection accuracy and in the cost of false positives.

3.5 Professional Knowledge and Research

AI changes the economics of regulatory and legal research in ways that compound across an organisation. A professional who previously spent two hours researching the JFSC's approach to a particular CDD question, or cross-referencing the Jersey AML/CFT/CPF Handbook against a client's structure, can now obtain a well-sourced analytical briefing in minutes, and redirect those two hours to work that genuinely requires their judgement.

This matters for the business model. The competitive advantage in professional services has historically accrued to firms and individuals with deep accumulated knowledge. AI democratises access to that knowledge base. The competitive advantage shifts toward those who can apply judgement effectively and who understand how to use AI as a force multiplier rather than a threat.

3.6 From Assistants to Agents


The most important shift in the AI landscape since early 2026 is the move from assistive tools to agentic systems, AI that does not merely draft a document on request but orchestrates an entire workflow: an onboarding agent that requests missing documents, extracts and verifies their contents, runs screening, assembles the CDD file and presents the completed package

with exceptions flagged for human decision. The JFSC's own Head of Innovation has described such systems as “digital colleagues, not just tools.”¹⁰

This is no longer speculative. The 2026 CCAF global survey found 52% of financial services firms actively adopting agentic AI, with 23% at the scaling or transforming stage, and fintechs are ahead of traditional institutions by 57% to 45%.¹¹ For fund administrators and CSPs, agentic systems are where the five operational areas above stop being five separate point solutions and become one connected operating model. The JFSC's July 2026 guidance addresses these systems by name, recommending stronger controls where AI can act without real-time human approval (Section 6): a governance rulebook for digital colleagues, not a prohibition of them. They are also where governance matters most, the more autonomous the system, the more deliberate the design of its human checkpoints must be. Section 7 returns to this.

4. The AI-Augmented Operating Model

The operating model that emerges from AI adoption looks meaningfully different from today's. It is not, in the near term, a model with dramatically fewer people. It is a model where the professional workforce is doing different and more valuable things. An AI capability layer sits beneath the professionals, absorbing high-volume procedural work with human review retained over every regulated output. That single structural change supports two distinct strategies: defend the margin (deliver the same output with a leaner team) or fund the growth (scale without hiring in proportion).



¹⁰Jersey Financial Services Commission, 'Agentic AI and what it means for Jersey's finance and legal sector: insights from Clairu Converge 2025'. Link in the Appendix.

¹¹2026 Global AI in Financial Services Report: see note 1.

Figure 1: The AI-augmented operating model. The AI capability embeds as a new layer beneath professional judgement, supporting either margin defence (leaner team) or growth without proportional hiring. The -30% shown in Panel A corresponds to the stretch case modelled in Section 4.1; the base case is approximately -23%.

Junior professionals today spend a significant proportion of their time on data extraction, document review, form completion and routine correspondence. In the AI-augmented model, these tasks are handled or drafted by AI with professional oversight. Junior professionals instead spend more time on review, quality control, client-facing communication and developing their judgement. These are activities that are better for their development, better for clients and more defensible under regulatory scrutiny.

Senior professionals and relationship managers are freed from operational firefighting to spend more time on client advisory work which is the highest margin activity in these businesses, and the activity most resistant to automation. A partner who currently spends 30% of their time on administrative coordination can redirect that capacity to business development, complex problem-solving or the kind of proactive client engagement that generates referrals and loyalty.

4.1 A Worked Example: The 40-Person Operations Function

Rather than assert headline percentages, this paper models them. Consider a 40-FTE operations function at a small to mid-size Jersey fund administrator or CSP: a representative mix of administrators, senior administrators, compliance staff, assistant managers and managers, costed at the midpoints of published Channel Islands salary bands. Its fully loaded cost is approximately £2.1 million a year.¹²

Team time is allocated across the five activity areas described in Section 2 plus a protected block (20% of all time) for judgement, advisory and client relationship work that is treated as entirely non-automatable. Within each activity, AI efficiency gains are applied only to the procedural, automatable share (50–75% of the activity depending on its nature) and adoption is phased over three years.

On these assumptions, the base case releases 22.8% of total team capacity at full adoption, the equivalent of just over nine FTE worth approximately £480,000 a year gross. Net of tooling licences, implementation and governance costs, the recurring benefit is approximately £417,000 a year, with around £127,000 net in year one and a cumulative net benefit approaching £870,000 over three years. A conservative case (materially lower efficiency gains on every activity) still releases 15% of capacity. A stretch case releases 30%, which is consistent with, and deliberately below, the best documented outcomes. Boston Consulting Group reports banks using agentic AI to cut KYC operating costs by up to 50%, with the strongest individual compliance and onboarding workflows showing efficiency gains of 60% or more.¹³

COST IMPACT A 40-person function releases the equivalent of 9 FTE in the base

¹²Alexander Daniels Offshore, 2025 Channel Islands Salary Guide, and companion Excel model: see note 4.

¹³Boston Consulting Group: 'The Know-Your-Customer Agentic AI Revolution' (2025) and 'A Faster Path to Scaling GenAI in Banking Compliance' (2025). These figures represent best-case outcomes for individual workflows at institutions with specialist AI teams, not firm-wide averages. Links in the Appendix.

COST IMPACT *A 40-person function releases the equivalent of 9 FTE in the base case: approximately £417,000 net per year, recurring. Scaled to a firm generating £30m of revenue, capacity release of this order is transformative. The full assumption set and sensitivity analysis are available in the companion Excel model.*

The wider evidence is consistent with this picture. The 2026 Global AI in Financial Services Report found that 75% of financial services firms reported positive AI productivity impacts in back-office and operations functions (the heartland of fund administration and corporate services) and that process automation is the single most widely adopted AI use case globally, at 79%.¹⁴ The financial gains are real where firms commit. Boston Consulting Group reports a global asset manager that reduced operating expenses by roughly a third through AI automation, delivering around US\$100 million of bottom-line impact,¹⁵ while PwC cites an institution that cut the cost of verifying commercial banking clients by 40% through AI-driven onboarding.¹⁶ These are outcomes already being measured at firms that moved earlier than their peers.

One cost line deserves emphasis rather than a footnote, integration. These estates run on established core platforms whose interfaces and data quality vary widely and in most real deployments the binding constraint is not model capability but data plumbing: extraction, cleansing and reconciliation before AI can touch the workflow. The worked example phases adoption over three years partly for this reason. Boards should expect integration and data remediation, not tooling licences, to dominate early implementation spend and should treat any proposal that arrives without an integration line as incomplete.

There are important nuances. The regulatory framework in jurisdictions like Jersey means that humans must remain in the loop for regulated activities. AI is not a substitute for the professional judgement that underlies regulatory advice, AML risk assessments or directorial decision-making. The appropriate model is 'human in the loop', AI handles volume, pattern, and procedure. Professionals handle judgement and accountability.

This is not a significant constraint. It is in fact a feature of the model that should be communicated proactively to regulators and clients alike. The intelligent firm is not one that has replaced its professionals with machines. It is one whose professionals are freed to bring their full capabilities to bear on the work that actually matters.

4.2 The Pricing Question: When the Hours Compress

There is a question the worked example raises that no efficiency model answers and it belongs on the same board agenda, much of this industry still bills time. Compress those hours under time-based billing and revenue falls before cost does, the efficiency gain is handed to the client by default. The recurring benefit modelled in Section 4.1 is captured only by firms whose pricing model moves with their operating model.

¹⁴2026 Global AI in Financial Services Report: see note 1.

¹⁵Boston Consulting Group cost-transformation analysis, reported via AI Magazine (2025). Link in the Appendix.

¹⁶PwC, 'How AI is reshaping banking'. Link in the Appendix.

The direction of travel is toward pricing outputs rather than effort, fixed fees per entity or per fund, service-level-based pricing and value pricing for advisory work. Firms with fixed fee books capture efficiency gains directly but firms billing hourly should treat AI adoption and pricing reform as a single programme, not two. There is also a defensive reading. As adoption spreads, competitive pressure will pass much of the savings into fees (the evidence on fintech-native entrants in Section 9 points exactly that way), so the end state for a late mover is not the old margin at a lower cost, it is the new market price carried on the old cost base. That more than any technology argument will be the true cost of standing still.

None of this is a reason to delay. It is a reason to put pricing on the transformation agenda from day one. The firms that redesign the operating model and the pricing model together will set the market's new price floor rather than be caught beneath it.

5. The AI-Native Firm: The Frontier Case

Everything to this point describes gains available to any well-governed firm within its existing architecture, it is what the Section 4 worked example models. This section describes something further, the case for firms prepared to rebuild around the technology itself. It is a frontier position, and we present it as such. The numbers here are not sourced from industry studies because, at the time of writing, we are not aware of any regulated firm who has walked this path. The numbers are our conviction about what is not only possible but, in our view, where this industry ultimately lands.




Figure 2: The three levels of AI maturity. Levels 1 and 2 optimise the firm you have; Level 3 rebuilds the firm around the technology.

5.1 From Renting Software to Growing It

The economics of software creation have collapsed. Prompt-native development platforms now allow working, production-quality internal systems to be built in weeks by the people who

understand the workflows best (the operations team members who live in them), working within guardrails set and owned by CTO and CISO functions. Real working MVPs can actually be produced in a matter of days. The leading platforms in this class now carry the attestations a regulated buyer expects (SOC 2 Type 2 compliance, ISO 27001:2022 certification, GDPR data-processing agreements with EU data residency and contractual commitments not to train on customer data) and support full code export. A system declared final can be lifted off the platform and into the firm's own estate entirely (sources in the Appendix). Private internal publishing on that footing consolidates a sprawling external SaaS estate onto a single certified, exportable dependency, reducing data sprawl and preserving an exit route that vendor SaaS never offers. One distinction must survive the enthusiasm, however, platform certification is not application security. What the firm builds on a certified platform still needs its own review gates, testing and periodic penetration testing, which is precisely what the guardrails described in this section and Section 5.3 exist to provide.

This inverts a decades-old assumption. Firms historically bought software because building cost more. For a large class of business applications, that premise no longer holds. In our experience a small to mid-size fund administrator carries an annual third-party software bill that runs well into six or even seven figures, across workflow tools, portals, trackers, reporting layers and internal utilities, much of it per-user pricing for products in which most features go unused. The Level 3 firm sets a deliberate roadmap. Build the replacement internally, run it in parallel, retire the contract. Applied across the long tail of the SaaS estate, the savings are substantial and fast. The honest counterweight: what the firm builds, the firm maintains. Retiring a contract retires the fee, not the lifecycle cost (support, patching and evolution move in-house), and the build-versus-buy decision should be made on that whole-life basis, contract by contract.

There is a second-order gain that never appears in a licence comparison, the workaround tax. Every firm pays it. The human glue between systems that almost fit, the re-keying, the export-to-spreadsheet rituals, the process contortions imposed by someone else's product design. Bespoke systems built around the firm's actual workflow eliminate the workaround, not merely the fee. And when the workflow changes, the system changes with it. Evolution on request within a governed change framework, rather than a feature request lodged with a vendor whose roadmap serves a thousand other customers.

5.2 What Cannot Be Built and How Much of It Can Be Automated Anyway

Honesty about the boundary strengthens the case. A firm cannot prompt its way to a global sanctions and PEP database: curated data of that kind is the genuine moat of the incumbent screening providers, and core regulated recordkeeping platforms similarly remain buy-side. However, the boundary sits much further out than most assume, because in most of these cost lines the data was never where the money went, the workflow was.

Sanctions screening is the worked example. Open-source sanctions datasets now offer commercially licensable coverage of the major regimes. An internally built AI screening agent can run continuous checks against such a dataset, contextually assess every hit, discard the false positives with reasoning attached and surface only genuine matches for human review. The analyst hours historically consumed clearing noise (the majority of screening cost) are removed and the premium per-user screening licence becomes negotiable or dispensable. The firm holds its hands up about what AI cannot do and in the same breath demonstrates how much of the process it transforms anyway. However, this is the part of the frontier case

that engages an MLCO's personal accountability, and it must be presented that way. Before any premium licence is allowed to lapse, four preconditions need evidencing, not asserting: benchmarked coverage against the incumbent provider, including ownership-based exposure, where the 50% rule demands curated ownership data that list-matching alone cannot supply; a defined update cadence with monitoring against designation changes; an extended parallel run, with discrepancy analysis reviewed and signed off by the MLCO; and dialogue with the regulator before, not after, the switch. For the officer whose name stands behind the control, data provenance, update cadence and coverage depth are not governance line items, they are the entire question, and the honest position is that most open datasets cannot yet evidence all four across every regime. Treat this worked example as a direction of travel with defined preconditions, not a procurement recommendation.

5.3 The People Model: Operators Become Builders

The Level 3 firm does not simply need fewer people, it needs different people, and most of them are already on the payroll. The administrator who knows exactly why the quarterly reporting process takes nine days is the ideal person to build the agent that does it in one, given the platform, the guardrails and the mandate. Staff redeploy from processing into AI development and agent operations, building, supervising, improving and auditing the automated estate. The client-facing team becomes a fraction of today's size and paradoxically more available to clients, because their time is no longer consumed by the machinery behind them. Instead, the machinery is working for them.

It is worth being concrete about the roles this firm hires and develops, because preparation starts now, and because almost none of them requires a computer science degree. AI agent builders, drawn from today's administrators and company secretaries, compose and configure agents on governed platforms. The scarce input is workflow expertise not software engineering. AI agent operators and supervisors run the automated estate day to day, monitoring queues, clearing exceptions and escalating what genuinely needs judgement. AI quality and audit analysts, a natural evolution for compliance and risk staff, own the sampling regimes, error metrics and evidence chains that keep human review substantive. A small professional engineering function, owned by the CTO/CISO, maintains the platform guardrails and any exported estate, and trainee agent operators are the new apprenticeship, the entry route through which the next generation learns the business.

That last role answers a question every board should ask. If AI absorbs the procedural work juniors have always learned the trade on, where do the senior professionals of 2032 come from? The answer has to be designed, not assumed. Judgment is built by doing, so the trainee pathway should rotate deliberately through manual execution of the core processes, periodically performing what the agents normally do, alongside supervised review, so that the skill of evaluating the machine's work rests on the experience of having done the work.

This is also where the redeployment framing of Section 4 and the leaner firm of this section reconcile, and the reconciliation should be stated plainly rather than blurred. At Level 2, released capacity redeploys into business development, advisory depth and client engagement. Level 3, pursued over years, will produce a smaller firm. The credible route to it runs through natural attrition, growth absorbing capacity and retraining into the roles above.

That is a multi-year workforce transition to be planned and communicated as such, not a restructuring dressed in softer language.




Figure 3: The intelligent firm — the future structure. New roles filled overwhelmingly from the existing workforce; the MLCO/MLRO accountability spine unchanged.

GOVERNANCE “Systems evolution on request” is only an asset inside a disciplined framework: CTO/CISO-owned development guardrails, change control, testing and rollback, access management, and audit trails on every deployment. Internal build changes the JFSC Outsourcing Policy analysis rather than removing it (the development platform, model APIs and cloud infrastructure beneath the estate remain third-party arrangements to assess on the Policy’s own terms), and it adds operational-resilience obligations on top. That trade should be made consciously and documented, not discovered. Section 7 sets out the wider risk framework.

5.4 The Prize, Stated as Conviction

Combining the effects (the retired SaaS estate, the eliminated workaround tax, the automated screening and processing workflows, and the redeployment of the workforce from processing to building), we believe headcount-equivalent gains of 50% and beyond are achievable, whichever strategic route the firm chose at Level 2, a materially leaner cost base, or growth at a scale the headcount-bound model could never fund. We state this plainly as our belief rather than dressing it as research because the evidence base does not yet exist, no regulated administrator has published the results of walking this path. The firms that walk it first will own that evidence, and the advantage it represents.



Figure 4: The work finds you – regulated sign-off arrives wherever the professional is, with authentication and a complete audit trail.

A working day, 2030

You leave the house without a bag. Your AI has been talking to you since the kettle boiled: the overnight emails triaged and half of them drafted in your voice awaiting your nod, the two board packs summarised with the one anomaly flagged, the screening run from 6am: forty-one hits, forty cleared with reasoning attached, one genuine match waiting for your judgement. You deal with it on the walk. By the time you reach the office (because you chose to come in today; a client is visiting), your morning admin is done and you haven't touched a keyboard.

There is no desk with your name on it. There are rooms for people, and glass where you need it: the laptop, the tablet, the lenses that are no longer a novelty. When something requires your signature rather than your delegation (the minutes, the CDD assessment, the NAV pack), it finds you, and you give it the one thing the machine cannot: genuine attention, and your name against the outcome. Authentication is your face. The audit trail is total. The work happens wherever you are; the accountability happens wherever you sign.

Your team is a fraction of its former size and better paid. The people who once cleared false positives now build and run the agents that clear them. The systems evolved last month because someone asked them to: no vendor roadmap, no change-request queue, no renewal negotiation. Your clients notice one thing above all: the humans they deal with have time for them.

The desk disappears. The sign-off doesn't.

6. The Regulatory Landscape: Jersey's Enabling Framework

For firms considering AI adoption in a regulated financial services context, the natural starting point is the regulatory framework. In Jersey, that picture is considerably more positive than many appreciate. The JFSC is not a reluctant observer of AI development. It is an active participant. Its public position, confirmed at the highest levels of the organisation in 2025 and 2026, is that responsible AI adoption is something it wants to support, not inhibit. On 14 July 2026 that position became formal as the Commission published its Guidance on the use of AI in Jersey's financial services sector, encouraging firms to explore AI within the obligations they already have.

6.1 The JFSC's Guidance on the Use of AI

The direction was signalled on 1 May 2026, when David Eacott, the JFSC's Executive Director of Supervision, published a blog setting out the regulator's approach following the Global Financial Innovation Network (GFIN) annual meeting, hosted that year in Jersey. The message was unambiguous: the JFSC is taking a "pragmatic and proportionate approach that supports innovation and growth, while continuing to safeguard customers and trust in financial services."¹⁷ Mr Eacott confirmed that dedicated guidance on artificial intelligence would follow shortly, and it has. On 14 July 2026 the Commission published its Guidance on the use of AI in Jersey's financial services sector, built around five principles for responsible AI use.¹⁸

The guidance is deliberately enabling. It creates no new requirements. It explains how existing legal and regulatory obligations apply when a firm uses AI, on a risk-based and proportionate footing. Its opening message is the argument of this paper made regulatory policy, firms are encouraged to explore AI, which 'can improve efficiency, automate routine tasks and increase productivity', while remaining accountable for outcomes just as they are with any other technology. Low-impact productivity tools need only light-touch controls, the expectations concentrate on higher-impact use cases such as customer onboarding, lending decisions and investment advice. Material AI use cases and their controls should, where appropriate, feature in the firm's business risk assessment.

Three features deserve a board's particular attention. First, the five principles (governance and accountability; cybersecurity, privacy and data quality; regulatory compliance; consumer protection and fair treatment; and transparency and explainability) are unpacked in Section 6.5. Second, the guidance addresses agentic AI directly by recommending stronger controls where systems can take actions, initiate workflows or carry out recommendations without real-time human approval. Third, five illustrative scenarios (AML onboarding, AI-assisted regulatory reporting, a third-party service outage, model drift in fraud detection and a biased lending model) show the principles applied in practice and each maps onto a control this paper

¹⁷David Eacott, Executive Director of Supervision, 'AI in financial services: enabling innovation with confidence', Jersey Financial Services Commission, 1 May 2026. Link in the Appendix.

¹⁸Jersey Financial Services Commission, Guidance on the use of AI in Jersey's financial services sector, issued 14 July 2026. Link in the Appendix.

recommends. Firms that establish a dialogue with the JFSC through the Innovation Hub and normal supervisory channels remain best positioned to demonstrate the compliance-first approach.

JFSC: 1 MAY 2026 *David Eacott, JFSC Executive Director of Supervision: “Supporting responsible adoption is a priority for the Jersey Financial Services Commission. In line with Jersey’s competitiveness focus, we are keen to work alongside firms as they explore new uses of AI. We particularly welcome early conversations where businesses are developing or scaling new approaches. We will publish our guidance on artificial intelligence shortly.”*

The JFSC’s Innovation Hub has been equally proactive. At Clairo Converge 2025, Andrew Evans, the Commission’s Head of Innovation, addressed agentic AI directly, framing such systems as digital colleagues capable of managing compliance tasks, onboarding workflows, and real-time regulatory updates with minimal human intervention: a significant opportunity for Jersey’s finance and legal sectors, provided it is embraced responsibly.¹⁹ Mr Evans was clear that AI does not change the accountability framework, firms remain responsible for the decisions they make, even where AI has assisted in making them and AI outputs deployed in critical functions should be sufficiently explainable.

6.2 AI Embedded in the JFSC’s Five-Year Strategy

The JFSC’s 2026–2030 strategy, published in March 2026, makes the Commission’s direction of travel explicit. Aligned with the Government of Jersey’s Competitiveness Programme, one of its four strategic pillars is “Support Growth”, defined as collaborating with government and industry, modernising the regulatory framework, and promoting innovation. On AI specifically, the strategy commits the Commission to “encouraging and enabling responsible innovation and technology adoption, including AI, across our regulated community and the JFSC,” adding that it recognises “the importance of leading by example here.” The strategy is unambiguous that AI, virtual assets and other digital innovations “present increased opportunities which competitive finance centres must continue to embrace.”²⁰

This is not regulatory boilerplate. It represents a five-year commitment from the Commission’s board to treat AI adoption as part of its core mission. For firms that have treated the regulatory stance on AI as uncertain, the 2026–2030 strategy removes that uncertainty.

6.3 The JFSC Leads by Example

The most persuasive evidence of the Commission’s genuine commitment to AI is that it has adopted the technology itself. In October 2024, the JFSC launched Reggie, an AI-powered regulatory chatbot drawing on the AML/CFT/CPF Handbook and Codes of Practice. By May 2025, Reggie had been upgraded to Reggie 2.0 with a knowledge base expanded by a further 48 documents, including laws, regulations, orders, official notices and guidance notes.

¹⁹JFSC, Clairo Converge 2025 insights: see note 10.

²⁰JFSC, 2026–2030 Strategy: see note 6.

This is not a one-off. The Commission's own 2026–2027 business plan commits it to “using AI to improve user experiences, by enabling more straight-through processing of administrative tasks and streamlining the examinations process.”²¹ In other words, the regulator is applying AI to precisely the categories of work (administrative processing and examination) that dominate the cost base of the firms it supervises. A regulator that deploys AI in its own operations and commits in its strategy to leading by example, understands the technology at first hand and is well placed to supervise its use by others. For regulated firms, that is a meaningful signal, the direction of travel is shared.

6.4 International Convergence: The FSB's Sound Practices

Jersey's posture sits within a converging international consensus. On 10 June 2026, the Financial Stability Board published a consultation on Sound Practices for the Responsible Adoption of AI by financial institutions: twelve practices structured around organisation-wide AI governance; risk management across the AI lifecycle; and cyber, ICT and third-party risk.²² The practices cover board-level oversight and risk appetite, data governance, explainability, human oversight, and the particular challenges of generative and agentic AI, and they map closely onto the governance framework recommended in Section 8 of this paper. Two practical points follow. First, a firm that builds its AI governance on these lines is preparing once for both the JFSC's guidance and the emerging global standard, a convergence the JFSC confirms itself, stating that its guidance closely aligns with the OECD AI Principles and guidance issued by international peer regulators. Second, the consultation window closes on 22 July 2026 with the final report expected in October 2026. Firms reading in time have a rare opportunity to shape how these practices apply to fiduciary and administration businesses, those reading later should treat the final report as the reference point, the preparation it demands is the same either way.

6.5 What the JFSC Expects: The Five Principles

The guidance rests on five linked principles, applied with regard to the risk and impact of each use case: governance and accountability; cybersecurity, privacy and data quality; regulatory compliance; consumer protection and fair treatment; and transparency and explainability. Firms should be designing their AI governance frameworks around them now. Five points of practical consequence for fiduciary and administration businesses follow.

Accountability stays with registered persons

The AML/CFT/CPF and regulatory obligations of Jersey-regulated businesses do not change because AI assists in performing them. The board, MLCO, MLRO and other responsible officers remain accountable for the quality of their firm's work, regardless of AI involvement, a point given fresh force by the revised AML/CFT/CPF Handbook that took effect on 31 May 2026 and the further MLCO-related changes that followed from July. This is the existing framework applied to a new context. Firms that build AI governance around professional review and sign-off (human-in-the-loop by design) will have no difficulty meeting this standard.

²¹Jersey Financial Services Commission, 2026–2027 Business Plan, published March 2026. Link in the Appendix.

²²Financial Stability Board, Sound Practices for the Responsible Adoption of Artificial Intelligence, consultation report, 10 June 2026. Responses close 22 July 2026; the final report is expected in October 2026. Link in the Appendix.

The guidance says as much in a single sentence: firms remain accountable for outcomes when they use AI, just as they do when they use any other technology.

Proportionality: the materiality ladder and the AI register

The most operationally useful element of the guidance is its materiality ladder: a structured guide to classifying each AI use case as low, medium or higher impact by reference to customer impact, regulatory reporting, financial crime relevance, data sensitivity, autonomy, complexity, reversibility and third-party reliance, with a proportionate governance response at each level, from a register entry and basic checks at the low end to senior approval, independent validation and explicit fallback at the top. Alongside it sits the AI register, a central record of material use cases capturing what the system is used for, the type of AI in plain terms, the main data inputs, who is responsible for oversight and a risk classification. Neither is a new requirement. Both are exactly the artefacts a supervisor will ask to see and a firm that adopts the governance framework in Section 8 will generate them as a by-product.

Transparency and explainability

The guidance asks firms to be able to explain AI-assisted decisions and outcomes ‘to a reasonable degree commensurate with risk and complexity’. A firm does not need to understand every technical detail of a model, but it should be able to give a clear account of the main factors that drove an outcome. This has practical implications for tool selection. Firms need to be able to describe, in straightforward terms, how an AI-generated CDD assessment or risk rating was reached and to evidence that a qualified professional reviewed it. Black-box AI producing regulated outputs without meaningful human oversight will not satisfy the guidance’s expectations.

The Outsourcing Policy

The guidance confirms that the current Outsourcing Policy, in force since January 2024, applies to outsourced activity that uses AI. It creates no additional requirements but applies the existing expectations, supplier due diligence, contractual oversight and incident notification, scrutiny of sub-contractors and cross-border transfers, and fallback or exit plans, proportionate to the materiality of the use case. Where an AI system performs a material part of a regulated activity (initial CDD assessments, regulatory filing preparation, AML monitoring), it is likely to constitute outsourcing within scope of the Policy and notification to the JFSC may be required. This is not a barrier to adoption, it is a practical compliance step that should be addressed at procurement stage, not retrospectively.

Internally built systems shift this analysis rather than escape it. The build itself may not constitute an outsourcing, but the model APIs, cloud infrastructure and development platform beneath it are third-party arrangements in their own right, to be assessed on the same terms, with operational resilience obligations added on top.

Data protection under the JDPL

AML and CDD data is sensitive personal data regulated by the Data Protection (Jersey) Law 2018. Jersey’s data protection framework holds a European Commission adequacy decision, providing internationally familiar protections for clients and counterparties. Firms deploying AI against client files, particularly where a third-party provider processes data outside the Crown Dependencies, must ensure their arrangements satisfy the JDPL. The guidance spells out what that means for AI, a recorded lawful basis for processing, impact assessments where

personally identifiable information is handled, data minimisation, lawfully managed cross-border transfers and updated privacy notices where AI materially changes how personal data are processed. This is a procurement stage question, not an afterthought.

6.6 Engaging with the JFSC's Innovation Hub

The JFSC's Innovation Hub exists specifically to support firms navigating new technology adoption. Its RegTech function provides regulatory guidance and actively reduces the barriers to adoption, including engaging directly with businesses developing or scaling new AI approaches. Early engagement with the Innovation Hub is sound risk management and an opportunity to shape how the new guidance is applied across the industry. The JFSC's GFIN membership, whose 2026 annual meeting it hosted, reinforces that the jurisdiction's regulatory posture is internationally aligned and forward-facing.

JERSEY FINANCIAL INCENTIVE *Jersey's 2024 Budget introduced a 150% super-deduction on qualifying expenditure on regulatory compliance technology (RegTech) for financial services companies. AI tools deployed for AML/CFT compliance, risk management, fraud prevention, and JFSC-required activities are likely to qualify. This tax advantage significantly improves the investment case for AI adoption and should feature in every internal business case presented to board or shareholders. (The relief was introduced as a pilot from the 2024 year of assessment: confirm its availability for current-year expenditure with Revenue Jersey or a tax adviser as part of the business case.)*

7. What Can Go Wrong: Managing the Risks

A paper that promised only upside would deserve scepticism. AI introduces genuine risks into a regulated operating model and the firms that adopt it credibly will be those that can articulate (to their boards, their clients and their regulator) exactly what can go wrong and how their governance catches it. Seven risks matter most.

Accuracy and hallucination

Large language models can produce confident, fluent output that is wrong. In a fiduciary context the discipline that answers this risk is simple to state: figures, dates, bank details and legal terms must always come from a verified source or a responsible member of staff. The AI drafts and reasons but never originates the numbers. The JFSC's guidance endorses that discipline in terms: its regulatory reporting scenario approves AI-assisted drafting precisely because figures and factual statements come from approved internal source data, with staff trained to identify fabricated or unsupported content. Combined with mandatory professional review before any output reaches a client, regulator or register, this converts hallucination from an existential objection into a managed quality-control problem. One refinement matters as volumes grow; the review must stay substantive. Automation bias is well documented (reviewers tend to approve what the machine produces), so the framework should evidence the review, not merely require it. Sampling-based quality assurance with tracked error rates, dual sign-off on high-risk outputs and monitoring of review time all help to ensure that a thirty-

second glance cannot masquerade as professional scrutiny. The supervisory question will not be whether a human was in the loop, but whether the firm can prove the loop did anything.

Data protection and shadow AI

The greater data risk in most firms is not the governed deployment but the ungoverned one: staff pasting client information into free consumer tools because the firm has provided no sanctioned alternative. The answer is twofold: an explicit AI use policy that names permitted tools, and enterprise-grade deployments in which client data is processed in appropriate jurisdictions under contractual data-processing terms, with techniques such as PII redaction applied before content reaches a model. Providing a good, sanctioned tool is the single most effective control against shadow AI.

Prompt injection and agent containment

The most consequential attack surface in agentic deployments is not the model but its inputs. An agent that reads inbound client documents, emails or side letters is reading untrusted content and adversarial instructions embedded in that content can attempt to redirect the agent's behaviour. Prompt injection sits at the top of the industry's LLM risk rankings for good reason and the answer is containment rather than trust: agents run with the least-privilege access required to tools and data; content from outside the firm is treated as data, never as instructions; outputs are validated before anything acts on them; and every material action an agent takes, not just the text it produces, is logged and attributable. The blast radius of a compromised agent should be a design parameter decided before deployment, not discovered after it.

Suspicion data and tipping-off

One category of data needs a harder rule than the general protocol: material connected to internal suspicious activity reports. Routing SAR-adjacent content through any external processor, however well-contracted, raises confidentiality and tipping-off considerations under the Proceeds of Crime framework that generic data-protection clauses do not answer. The data-handling protocol should say so explicitly, files subject to an internal disclosure are flagged and excluded from AI processing pipelines unless the MLRO has approved the specific use and the exclusion itself is evidenced. Few firms' AI policies will address this today, those that do will find it a differentiator in any supervisory conversation.

Vendor concentration and model availability

A firm whose onboarding workflow depends on a single AI model has acquired a new operational dependency and June 2026 provided a vivid illustration. The most capable frontier model on the market was suspended by its own developer for 19 days, weeks after launch, when researchers demonstrated a jailbreak, before being redeployed with strengthened safeguards.²³ A well-run vendor pausing a product on safety grounds should be the system working, not failing. But the operational lesson stands, model availability is now a resilience issue in exactly the sense the JFSC's outsourcing framework contemplates. Firms should know which of their workflows depend on which models and have a documented fallback (an alternative model, or a manual procedure) for each material dependency. The guidance's third illustrative scenario makes the same point, when a third-party AI service fails, the well-governed firm pauses or reviews the affected workflow rather than letting it run unchecked.

²³See note 5.

The same discipline applies with greater force to the internally built estate described in Section 5.

Explainability and auditability

Section 6 sets out the JFSC's expectation that outputs in critical functions be sufficiently explainable. Operationally this means retaining the audit chain, what was the input, which model and version produced the output, what did the reviewing professional change and who signed it off. Tools should be selected partly on their ability to evidence this chain, a system that cannot show its working will eventually fail a supervisory conversation.

Liability and insurance

Professional indemnity policies were not drafted with AI-assisted work in mind. Firms should confirm with brokers how their cover responds to errors in output assisted by AI and ensure the human-review framework is documented well enough to demonstrate that professional judgement (the risk the policy insures) remained in the loop.

One further consideration for firms with EU-facing operations, the EU AI Act's obligations are phasing in through August 2026 and administrators operating in or servicing structures in Luxembourg or Dublin should map whether any of their AI use falls within its scope. Jersey firms are outside the Act's direct reach for domestic activity, however, the Act will shape the expectations of EU counterparties and clients.

None of these risks are a reason not to adopt. Each is a reason to adopt deliberately, with a governance framework that makes the firm's answer to "what if it goes wrong?" a strength of the proposition rather than a gap in it.

8. A Practical Roadmap for Adoption

The firms that will benefit most from AI are not those that launch the most ambitious transformation programmes. They are those that identify the highest-value and most implementable use cases, prove them quickly, build internal confidence and expand from there. Transformation as a slogan is a distraction. Transformation as a disciplined operational programme, with clear milestones and measurable outcomes, is a competitive advantage.

A sensible adoption roadmap proceeds in three phases:

Phase	Focus	Key actions
Phase 1 Foundation (Months 1–6)	Audit, select and govern	Map current processes by volume and time cost. Evaluate AI tooling with attention to data security and regulatory compliance. Select 2 or 3 high-value pilot use cases. Establish AI policy, AI register and human-review governance.
Phase 2 Build (Months 6–18)	Deploy, measure and expand	Deploy pilots and measure rigorously. Expand to adjacent processes based on evidence. Integrate AI tooling into core operational platforms. Build internal AI literacy across the

		professional workforce.
Phase 3 Optimise (Months 18+)	Differentiate and lead	Review operating model implications: workforce planning, pricing, client proposition. Decide the maturity question deliberately: remain at Level 2 or set a Level 3 roadmap (Section 5). Build proprietary capability where competitive advantage can be established. Contribute to industry and regulatory dialogue on standards.

Selecting the Right Pilot Use Cases

Not all AI use cases are created equal. The best initial candidates share three characteristics: they involve a high volume of structured or partially structured work; the quality of the output can be evaluated objectively by a professional reviewer; and the consequences of an error are recoverable. KYC document review, board minute drafting and regulatory deadline monitoring are strong candidates for most firms in this sector. The JFSC's guidance illustrates the first of these itself. Its lead scenario is an AI tool supporting AML onboarding, approved through governance, recorded in the AI register, tested against known red-flag cases and with higher-risk cases routed to a suitably skilled person.

In the first phase avoid use cases where errors carry high immediate consequences and where there is limited ability to verify outputs before they reach clients or regulators. Build confidence in the technology and the governance framework first.

Governance: The Non-Negotiable

An AI adoption programme without a governance framework is an operational and regulatory liability. Every firm deploying AI in a regulated context needs, at minimum, an AI use policy articulating what AI tools may and may not be used for, a data-handling protocol confirming how sensitive data is processed and stored, an AI register of material use cases with risk classifications, human review requirements specifying when and by whom AI outputs must be reviewed before use, a vendor and model-availability fallback plan and an incident response process for when AI tools fail or produce materially incorrect output. A framework built on these lines will map naturally onto the FSB's twelve sound practices and onto the five principles and materiality ladder of the JFSC's guidance on the use of AI.²⁴

This is not complex to design. It is, however, frequently overlooked in the enthusiasm of early deployment, with predictable consequences.

9. The Cost of Standing Still

The temptation for established firms is to wait, to let others experiment, manage the regulatory uncertainty from a safe distance and adopt proven tools once the landscape is clearer. This is understandable risk management. It is also, in all likelihood, a costly error.

The economics of AI adoption create compounding advantages for early movers. The firm that compresses its cost base meaningfully in year one generates a margin improvement that

²⁴Financial Stability Board consultation: see note 21.

funds further investment in year two. The firm that builds proprietary AI capability in, say, Jersey regulatory compliance will have a tool that competitors cannot easily replicate, not because the underlying technology is scarce but because the training data, the institutional knowledge and the workflow integration take time to build. That is the honest mechanism behind first-mover advantage in this sector (accumulated workflow learning, retrained people and owned evidence, none of which can be bought quickly), and it deserves stating plainly, because the counterargument is respectable: deliberate fast-following has often been the right strategy in this industry, letting others absorb the tooling and regulatory risk. What has changed is what decays while waiting. Every quarter of delay is a quarter of workflow learning, staff capability and client evidence accruing to someone else, and waiting does not preserve today's price, because as adoption spreads competition passes much of the efficiency gain into fees (Section 4.2). The late mover meets the new market price with the old cost base.

The competitive landscape is also widening. The 2026 CCAF global survey found that only 6% of traditional financial institutions have reached the transforming stage of AI deployment, compared with 19% of fintech firms. 44% of incumbents remain at the piloting stage.²⁵ Fintech-native fund administration firms entering the market are not starting from piloting. They are building AI-first operating models from the outset with structural cost advantages that established players who have not adapted will find very difficult to match. This is not a threat on the horizon. It is the dynamic already shaping today's competitive market.

The talent dimension compounds the pressure. The next generation of professionals want to work in environments where they are doing interesting, high-value work, not grinding through manual KYC files or drafting the same board resolution for the hundredth time. Firms that invest in AI will find it meaningfully easier to attract and retain the talent that drives quality and client relationships. Those that do not will find attrition rising among exactly the people they most want to keep.

The clients of these businesses are also becoming more sophisticated about operational efficiency. A fund manager choosing between administrators increasingly asks not just about fees and service quality, but about how the provider operates, and what the operational model means for error rates, turnaround times and the quality of the human engagement they will actually receive. An AI augmented operating model is an increasingly compelling answer to that question.

10. Conclusion

The question facing leadership teams in fund administration and corporate services is not whether artificial intelligence will transform their industry. It will, and it is already beginning to. The question is whether their organisations will be in front of that transformation, shaping it to their advantage, or behind it responding to competitive pressure with diminishing options.

The operating model of the intelligent firm is achievable today, with existing tools. It requires no speculative technology bets and no wholesale destruction of the people and relationships that have built these businesses. It requires strategic clarity about where the highest-value use cases lie, governance discipline to deploy AI safely and credibly in a regulated environment, and the organisational will to challenge ways of working that have been accepted as fixed for too long.

²⁵2026 Global AI in Financial Services Report: see note 1.

The professionals who work in these businesses, whether in the Channel Islands, Dublin, Luxembourg or Cayman, are too well trained and too expensive to spend their days on work that a machine can do better. The firms that understand this and act on it will build businesses that are more efficient, more scalable, more resilient and more focused on the advisory work that clients genuinely value and that professionals genuinely want to do.

AI does not replace the judgement of the experienced professional. It removes the procedural burden that has always competed with that judgement for time and attention.

The window for first-mover advantage is open. It will not remain open indefinitely.

About Prospero Advisory

Prospero Advisory Limited is a Jersey-incorporated governance and operational advisory boutique (registered number 166208), working with fund administrators, corporate service providers, fund managers, investors and boards. Its practice spans three connected arms: governance and board advisory; operational advisory, including provider due diligence and fractional COO support; and AI adoption: the strategy, governance frameworks and practical implementation this paper describes.

Prospero's founder, Paul Monahan, brings some nineteen years in Jersey's regulated financial services industry, including senior executive roles at two of the sector's largest administrators and a period as Manager, Funds Authorisation at the Jersey Financial Services Commission, providing a first-hand perspective on how Jersey's regulatory framework operates in practice. He has held non-executive and chair roles across Jersey fund, GP, trustee and corporate structures for more than a decade and in his personal capacity is registered with the JFSC for specified Schedule 2 business under the Proceeds of Crime (Supervisory Bodies) (Jersey) Law 2008.

Prospero also builds what it recommends. Prospero Tools is the firm's AI product suite, designed specifically for the fiduciary and fund administration sector: **Brief**, a board-pack reviewer that produces multi-lens analysis of board materials before directors walk into the room, and **Minutes**, a minute-drafting tool built for trust company and fund administration teams. Both are deployed under the governance approach described in this paper: enterprise security, EU-resident AI processing, in-house PII redaction and the working principle that figures, dates and legal terms are never originated by the AI. The suite is currently in invitation-only beta.

This white paper is provided for informational purposes only. Nothing in this paper constitutes legal or regulatory advice and readers should seek specific professional guidance before taking any action in reliance on its contents.

Paul Monahan | Founder | Prospero Advisory Limited | paul@prosperoadvisory.je | +44 7797 922 174

prosperoadvisory.je | prosperotools.com | Jersey, Channel Islands

Appendix: Sources and References

All factual claims in this paper are drawn from the published sources below, each verified as at July 2026. Each entry notes the claims it supports. Links are clickable in the digital edition.

Adoption and market data

2026 Global AI in Financial Services Report. Cambridge Centre for Alternative Finance with the BIS, IMF and World Economic Forum, April 2026. Supports: 81% of firms adopting AI; 75% reporting positive back-office productivity impacts; process automation the most adopted use case (79%); 6% of incumbents vs 19% of fintechs at the transforming stage (44% of incumbents piloting); agentic AI actively adopted by 52% of firms, 23% at scaling/transforming, fintechs 57% vs incumbents 45%. [jbs.cam.ac.uk \(CCAF publications\)](https://jbs.cam.ac.uk/CCAF/publications)

Finastra, Financial Services State of the Nation Survey. 1,509 senior executives across 11 markets. Supports: only 2% of financial institutions report no use of AI. [finastra.com \(press release\)](https://finastra.com/press-release)

Fenergo, Financial Crime Industry Trends 2025. Supports: use of advanced AI tools in KYC/AML rising from 42% (2024) to 82% (2025), and periodic KYC review automation averaging roughly one-third. [resources.fenergo.com \(newsroom\)](https://resources.fenergo.com/newsroom)

Fenergo, cost of KYC compliance. Supports: average corporate KYC review cost of approximately US\$2,500–2,600. [resources.fenergo.com \(KYC compliance costs\)](https://resources.fenergo.com/KYC-compliance-costs)

Boston Consulting Group, The Know-Your-Customer Agentic AI Revolution (2025). Supports: banks using agentic AI to cut KYC operating costs by up to 50%. [bcg.com \(publication\)](https://bcg.com/publication)

Boston Consulting Group, A Faster Path to Scaling GenAI in Banking Compliance (2025). Supporting source for efficiency gains in AI-assisted compliance workflows. [bcg.com \(publication\)](https://bcg.com/publication)

Boston Consulting Group, cost-transformation analysis (via AI Magazine, 2025). Supports: a global asset manager reducing operating expenses by roughly one-third through AI automation, for approximately US\$100 million of bottom-line impact. [aimagazine.com \(article\)](https://aimagazine.com/article)

PwC, How AI is reshaping banking. Supports: a 40% reduction in the cost of verifying commercial banking clients through AI-driven onboarding and verification. [pwc.com \(financial services library\)](https://pwc.com/financial-services-library)

Jersey regulatory sources

JFSC, Guidance on the use of AI in Jersey's financial services sector. Issued 14 July 2026. Supports: the five principles for responsible AI use; the confirmation that the guidance creates no new requirements and is risk-based and proportionate; the materiality ladder and AI register; the agentic AI controls; the five illustrative scenarios; and the stated alignment with the OECD AI Principles. [jerseyfsc.org \(news\)](https://jerseyfsc.org/news)

JFSC, 'AI in financial services: enabling innovation with confidence'. Blog by David Eacott, Executive Director of Supervision, 1 May 2026. Supports: the quoted regulatory

posture and the commitment to publish AI guidance, delivered 14 July 2026. [jerseyfsc.org \(industry update\)](https://jerseyfsc.org/industry-update)

JFSC, 2026–2030 Strategy. Published March 2026. Supports: the ‘Support Growth’ pillar; the verbatim commitment to responsible innovation and technology adoption ‘including AI, across our regulated community and the JFSC’ and to ‘leading by example’; and the quoted language on competitive finance centres. [jerseyfsc.org \(publications\)](https://jerseyfsc.org/publications)

JFSC, 2026–2027 Business Plan. Published March 2026. Supports: the commitment to ‘using AI to improve user experiences, by enabling more straight-through processing of administrative tasks and streamlining the examinations process’. [jerseyfsc.org \(news\)](https://jerseyfsc.org/news)

JFSC, Agentic AI insights from Clairo Converge 2025. Supports: Andrew Evans’s ‘digital colleagues’ framing and the accountability and explainability principles. [jerseyfsc.org \(news\)](https://jerseyfsc.org/news)

JFSC, ‘Meet Reggie in beta’ (October 2024). Supports: launch of the JFSC’s AI-powered regulatory chatbot. [jerseyfsc.org \(news\)](https://jerseyfsc.org/news)

JFSC, Reggie 2.0 (May 2025). Supports: expansion of Reggie’s knowledge base to over 48 documents. [jerseyfsc.org \(news\)](https://jerseyfsc.org/news)

JFSC, AML/CFT/CPF Handbooks. Supports: the revised Handbook effective 31 May 2026 and subsequent MLCO-related changes. [jerseyfsc.org \(financial crime\)](https://jerseyfsc.org/financial-crime)

JFSC, Outsourcing Policy. In force since 1 January 2024. Supports: the outsourcing analysis applied to material AI deployments. [jerseyfsc.org \(guidance and policy\)](https://jerseyfsc.org/guidance-and-policy)

Government of Jersey, RegTech super-deduction. Official guideline for the 150% super-deduction on qualifying RegTech expenditure, available from year of assessment 2024. [gov.je \(income tax guidelines\)](https://gov.je/income-tax-guidelines)

International regulatory sources

Financial Stability Board, Sound Practices for the Responsible Adoption of AI.

Consultation report, 10 June 2026; twelve sound practices across governance, lifecycle risk management, and cyber/ICT/third-party risk. Responses close 22 July 2026; final report expected October 2026. [fsb.org \(consultation\)](https://fsb.org/consultation)

Financial Stability Board, consultation report (full text). PDF of the twelve sound practices. [fsb.org \(report PDF\)](https://fsb.org/report-PDF)

European Commission, EU AI Act. Overview of the regulatory framework for AI and its phased application. digital-strategy.ec.europa.eu

Technology developments

Anthropic, Claude Fable 5 and Claude Mythos 5. Launch announcement, 9 June 2026. Supports: the mid-2026 frontier capability step, including long-horizon autonomous work. [anthropic.com \(announcement\)](https://anthropic.com/announcement)

Anthropic, Redeploying Claude Fable 5. 1 July 2026. Supports: the 19-day suspension and redeployment cited in Section 7 as a model-availability case study. [anthropic.com \(announcement\)](https://anthropic.com/announcement)

Lovable, Security & Trust Centre. Example of the prompt-native development platform class referenced in Section 5. Supports: SOC 2 Type 2 compliance and ISO 27001:2022 certification, GDPR data-processing agreement with EU data residency, no-training commitment, and full code export. trust.lovable.dev

Cost model inputs

Alexander Daniels Offshore, 2025 Channel Islands Salary Guide. Supports: all salary bands used in the Section 4 worked example (midpoints of published ranges for trust, fund, compliance and corporate services roles). alexanderdanielsoffshore.com ([salary guide PDF](#))

Companion model: The Intelligent Firm AI Cost Impact Model (Excel) contains the full assumption set, three-year phasing, and sensitivity analysis behind every figure in Section 4. Available from Prospero Advisory on request.

Regulatory instruments referenced: the JFSC Codes of Practice; the JFSC Outsourcing Policy (in force from January 2024); the Financial Services (Jersey) Law 1998; the Proceeds of Crime (Jersey) Law 1999 (including Schedule 2); and the Data Protection (Jersey) Law 2018, all available in full from jerseyfsc.org and gov.je.