MatchPoint Research Working Paper No. 2026/T07

Building the AI-Enabled Investment Firm: A Technology Architecture for Modern Advisory

A governed decision, evidence, workflow and resilience architecture for family offices and fund managers

Chennakeshav (CK) Adya
Independent Researcher
July 2026

ABSTRACT

Background. Investment firms operate through connected commercial, investment, control and management decisions. AI can increase analytical capacity while creating risks around evidence, authority, confidentiality, third parties and resilience. **Objective.** This paper develops a governed technology architecture for family offices, fund managers and modern advisory firms. **Approach.** The analysis reviews 19 official, supervisory, voluntary-framework and practitioner sources. Regulatory scope, consultation status and survey populations are retained. **Findings.** The Bank of England and FCA's 2024 survey of 118 responding firms reported current AI use by 75% of respondents, one-third third-party implementation across reported use cases and complete understanding of AI technologies by 34% of firms. The proposed architecture uses decision objects, authoritative systems of record, a permissioned evidence graph, deterministic financial services, registered models, workflow policy, human approval, zero-trust access and end-to-end observability. **Implications.** Firms can establish identity and evidence first, test assistive components, shadow complete workflows and expand into controlled production after task-level quality, authority, service, cost and incident evidence meets approved thresholds.

Highlights

- Decision objects connect evidence, calculations, models, challenge, authority and monitoring.
- Systems of record remain authoritative for defined business objects and approved state.
- Deterministic services, statistical models, language systems and agents receive separate controls.
- Identity, permissions, security, records and observability cross every architecture layer.
- Productivity claims require accepted packets, reviewer effort, service, cost and incident evidence.

JEL Classification: G23, G24, G34, L86, M15, O32, O33

Keywords: artificial intelligence, investment management, family office, fund management, enterprise architecture, data governance, workflow orchestration, model risk, operational resilience, human oversight

1. INTRODUCTION

An investment firm is a chain of decisions supported by evidence, models, judgement, communication and control. An opportunity enters the pipeline. Research and diligence shape a view. A model expresses economics and risk. An investment committee authorises a position. Portfolio teams monitor outcomes. Client, investor and regulatory reporting explain what happened and why. Technology touches every stage, yet many firms still operate through disconnected inboxes, shared drives, spreadsheets, data terminals, customer-relationship systems, portfolio platforms and presentation files.

Artificial intelligence can increase the analytical capacity of this operating model. Document models can classify and extract. Language models can retrieve, compare and draft. Statistical models can estimate risk or identify anomalies. Workflow agents can route tasks and use controlled tools. These capabilities create value only when the firm can preserve evidence, identity, permissions, calculation integrity, review and decision authority.

The institutional evidence shows rapid adoption alongside incomplete understanding. The Bank of England and Financial Conduct Authority's 2024 survey received 118 responses across banks, insurers, non-bank lenders, investment and capital-markets firms, and financial-market infrastructure and payments firms. Seventy-five per cent of respondents reported current AI use and another 10% planned use within three years. Foundation models represented 17% of reported use cases. One third of use cases were third-party implementations; the three most commonly named model providers represented 44% of named model providers. Forty-six per cent of firms reported only partial understanding of the AI technologies they used, compared with 34% reporting complete understanding [1]. These results are respondent-reported and cover financial services broadly.

The same survey reported some automated decision-making in 55% of AI use cases and full autonomy in only 2% [1]. The distinction between assistance and authority is therefore an operating-model question. A modern advisory firm needs an architecture that decides where data lives, which system is authoritative, which task a model may perform, how an output is tested, who reviews it, when a process stops and what record remains.

This paper develops a technology architecture for family-office investment teams, private-markets managers, fund managers, placement teams and related advisory firms. It covers origination, research, diligence, investment analysis, committee governance, portfolio monitoring, investor relations, reporting and enterprise control.

The central design proposition is that architecture makes the operating model executable. The architecture should expose the firm's mandate, decision rights, evidence, calculations, models, workflows and records as governed components. It should support people in making more complete and timely decisions while retaining accountable human authority over material conclusions, commitments and external communications.

The paper is a research synthesis and design framework. It does not report an observed Matchpoint production deployment, an observed productivity gain, realised revenue, investment performance or a comparison between specific technology vendors.

2. DEFINITIONS, SCOPE AND METHOD

2.1 Investment-firm scope

The term investment firm is used functionally. It includes an organisation that originates, evaluates, structures, approves, manages or advises on investments and capital relationships. The operating model may sit inside:

- a single-family or multi-family office;
- a private-equity, private-credit, real-estate or venture manager;
- an alternative-investment adviser;
- a placement or capital-advisory firm;
- an investment office within a holding company; or
- a specialist corporate-finance adviser.

Legal and regulatory status varies by entity, service, client, activity and jurisdiction. References to rules or supervisory material retain their stated scope. A family office, exempt manager or unregulated advisory activity should not be assumed to fall inside a regime merely because a similar operating control is useful.

2.2 Technology classes

This paper separates five technology classes.

Systems of record hold authoritative entities, relationships, transactions, positions, documents, accounting records, approvals and communications.

Deterministic services apply fixed rules or formulas. Examples include cash-flow calculations, exposure aggregation, fee calculations, covenant tests and reconciliations.

Statistical and machine-learning models estimate, rank, classify or detect patterns. They require defined intended use, development evidence, validation, monitoring and change control.

Language and document models classify documents, propose extracted fields, retrieve passages, compare clauses, summarise evidence and draft analysis. Their outputs require source and claim verification.

Workflow agents coordinate retrieval, tools, state transitions, exceptions and review. Their authority is established by policy, permissions and approval gates.

This separation supports component-specific controls. A deterministic calculation should be recomputed. A risk model should be tested for performance and stability. A generated paragraph should be checked for citation validity and unsupported claims. An agent should be tested for correct tool use, stop behaviour and authority compliance.

2.3 Evidence method

The paper reviews 18 primary, official, supervisory and practitioner sources.

Financial-sector AI evidence. Bank of England and FCA survey results, IOSCO reports, the FSB consultation and ESMA's statement describe use cases, adoption, third-party reliance, governance and investor-protection risks [1-5]. IOSCO's March 2025 report and the FSB's June 2026 report are consultations; they do not create binding international standards [3,4].

Conduct and jurisdictional evidence. SEC enforcement, CBUAE guidance, joint UAE supervisory guidance, DORA, FCA operational-resilience material, DIFC data-protection material and DFSA cyber expectations inform disclosure, data, resilience, outsourcing, records and governance [6-10,16-18]. Applicability depends on the firm and jurisdiction.

Cross-sector technology frameworks. NIST's AI Risk Management Framework, Generative AI Profile, Cybersecurity Framework 2.0, Zero Trust Architecture and Secure Software Development Framework provide voluntary, cross-sector risk and architecture references [11-15].

Practitioner architecture. OpenAI's guide contributes provider-authored agent design patterns [19]. It is not independent comparative evidence.

Evidence class	Permitted use	Boundary retained
Official survey	Adoption, autonomy, understanding and third-party context	Respondent-reported; financial services broadly
Supervisory or regulatory material	Conduct, governance, record, resilience and control objectives	Jurisdiction and entity scope retained
Consultation	Emerging risk taxonomy and proposed practices	Subject to consultation; not binding
Voluntary framework	Cross-sector architecture and risk-management outcomes	Requires firm-specific implementation
Provider guidance	Agent design and tool-orchestration patterns	Provider-authored; no independent performance claim
Matchpoint synthesis	Target architecture, operating model, scorecard and roadmap	Requires organisation-specific validation
Matchpoint scenario	Economic calculation structure	Unverified management assumptions only

3. THE ARCHITECTURE PROBLEM

3.1 Fragmented decision state

Investment work rarely occurs inside one application. A relationship may begin in a customer-relationship system. Documents arrive by email or enter a data room. Analysts maintain spreadsheets and models. Research sits in terminals, subscriptions and internal notes. Diligence issues appear in trackers. Committee materials become presentations and PDFs. Approved terms move into legal documents and portfolio systems. Reporting teams rebuild explanations for clients and investors.

Each hand-off can lose identity, version, source, definition or approval state. The newest file may not be the approved file. A conclusion can remain after its supporting assumption changes. A relationship owner may see a different pipeline status from the investment team. Portfolio monitoring may use a covenant definition that differs from underwriting. Investor reporting may quote a number whose calculation cannot be reproduced.

AI can accelerate these hand-offs and can also accelerate the propagation of a weak record. The target architecture therefore begins with stable identifiers, explicit state and traceability.

3.2 Four coupled systems

The firm contains four coupled systems.

- **The commercial system** manages relationships, opportunities, mandates, fundraising and communications.
- **The investment system** manages research, diligence, modelling, committee decisions, positions and monitoring.
- **The control system** manages permissions, compliance, conflicts, privacy, cyber risk, records, incidents and independent challenge.
- **The management system** manages people, capacity, economics, service quality and change.

A workflow may cross all four. A potential co-investment involves relationship evidence, investment analysis, conflict review, data permissions, committee authority, allocation records and investor communications. Architecture should preserve those relationships while limiting each user's and service's access to the minimum required.

3.3 The source of truth problem

A single universal database is rarely practical. A firm can define authoritative ownership by object instead.

Object	Example authoritative system	Required relationship
Person and organisation	CRM or master-data service	Links to funds, deals, communications and permissions
Opportunity and mandate	Deal or mandate platform	Links to owner, stage, scope, conflicts and approvals
Document and evidence	Governed document repository	Links to entity, version, permission, fact and decision
Calculation and model	Controlled analytics service or model registry	Links to inputs, code version, validation and output
Investment decision	Committee and approval register	Links to evidence, conditions, dissent and authority
Position and transaction	Portfolio or accounting system	Links to approved terms and monitoring obligations
Communication	Approved communication and archive system	Links to recipient, evidence, reviewer and record class

The integration layer should move identifiers and governed events rather than duplicate every object into every system. A relationship update, new document, changed model, committee approval or covenant breach becomes an event with identity, timestamp, owner and downstream consumers.

3.4 Authority and representation

The SEC's March 2024 orders found that two investment advisers made false and misleading statements about their use of AI and imposed USD 400,000 in total civil penalties [6]. The enforcement facts are specific to those firms. The architecture lesson is wider: public statements about technology need an evidence owner, approval path and retained basis.

The same discipline applies to investment conclusions. A model may propose a classification. A deal team may interpret it. An independent function may challenge it. A committee may decide. The system should represent each action distinctly.

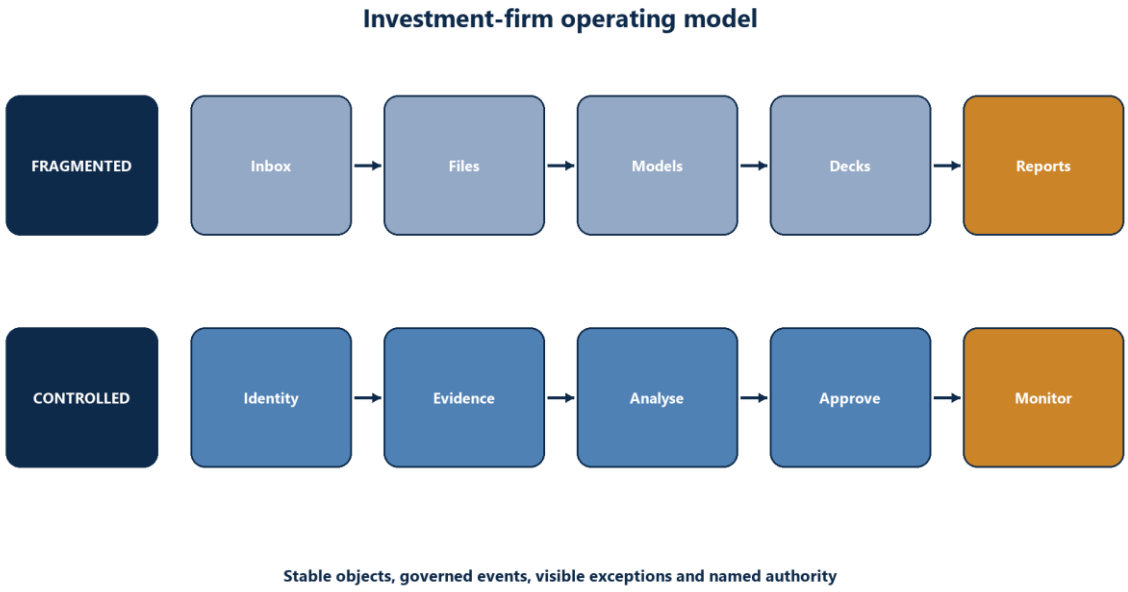

Figure 1. Fragmented investment-firm hand-offs compared with an event-driven evidence-to-decision operating model
Source: Matchpoint framework.

4. TARGET OPERATING MODEL

4.1 Decision objects

The controlled unit of work is a decision object. It can represent a qualification decision, diligence conclusion, valuation, allocation, committee approval, portfolio action or external statement.

A decision object records:

- objective and scope;
- entities and instruments;
- permitted evidence and cut-off;
- assumptions and adjustments;
- calculations and model outputs;
- materiality and risk class;
- exceptions and unresolved matters;

- reviewer and decision authority;
- conditions and monitoring consequences; and
- version and action history.

The object can have multiple work packets. A diligence decision object may include source registration, commercial analysis, financial analysis, technology analysis, legal issues and an integrated challenge packet. Each packet has its own tools, tests and owner.

4.2 Product-oriented platform teams

Technology ownership should align with business capabilities rather than isolated applications. A small firm may combine roles; the responsibilities remain explicit.

Platform capability	Business owner	Technology owner	Control partner
Relationship and origination	Head of origination	CRM or product lead	Compliance and conflicts
Research and evidence	CIO or research head	Data and knowledge lead	Information governance
Deal execution	Investment or advisory lead	Workflow product lead	Legal, risk and compliance
Portfolio intelligence	Portfolio head	Portfolio-data lead	Valuation and risk
Investor and client reporting	Investor-relations head	Reporting product lead	Compliance and records
Enterprise AI	Senior accountable executive	AI platform lead	Independent risk, privacy and security

Product teams maintain an outcome, user journey, data contract, service level, control set, backlog and cost. Application ownership remains relevant, though it sits within the end-to-end capability.

4.3 Three lines of accountability

The first line owns the use case and outcome. It defines the purpose, evidence, workflow, controls and acceptance criteria. The second line provides independent challenge across risk, compliance, legal, privacy, model risk, information security or valuation as applicable. Internal audit or another independent assurance function assesses design and operation according to the firm's governance model.

For a smaller family office, named external advisers may provide parts of the second or third line. The retained decision authority and escalation path should be documented.

4.4 A common service catalogue

Shared services reduce duplicated controls. The architecture should offer reusable capabilities for:

- identity and access;
- entity resolution;
- document registration;
- permission-aware retrieval;
- structured extraction;
- deterministic calculation;
- model registration and serving;
- prompt and tool policy;
- workflow state;
- human approval;
- communication release;

- immutable logging;
- evaluation and monitoring; and
- incident and exception management.

The catalogue separates common infrastructure from investment judgement. A sector team can build a specialist underwriting packet on common source, identity, calculation and approval services.

5. REFERENCE TECHNOLOGY ARCHITECTURE

5.1 Experience layer

Users interact through role-specific workbenches. An originator sees relationships, pipeline, mandate fit and next actions. An analyst sees sources, models, open questions and draft work. A committee member sees decision-ready evidence, changes, exceptions and approvals. A portfolio manager sees positions, assumptions, covenants and monitoring events. Investor relations sees approved messages, fund data and communication status.

The experience layer may include chat or natural-language search. Every result should retain its source, permission and status. A conversational interface should not create a separate ungoverned memory of the firm.

5.2 Workflow and policy layer

The workflow engine holds process state, dependencies, due dates, approvals, materiality, tool permissions and stop conditions. Policy is expressed as rules that can be tested.

Examples include:

- a conflict check must pass before external engagement;
- an unverified fact cannot support a released investment conclusion;
- a financial ratio must use an approved definition and calculation version;
- a committee pack cannot be released while a material exception is unresolved;
- an external recipient must match an approved entity and communication purpose; and
- a portfolio breach must route to the named authority within the approved window.

The policy layer creates exceptions with an owner. It should not silently fill a missing approval.

5.3 Intelligence layer

The intelligence layer contains retrieval, document models, language models, statistical models, rules and deterministic services. A router selects the appropriate component according to task and risk.

Evidence retrieval returns approved passages and structured facts. Deterministic tools perform calculations. Statistical models estimate bounded outputs. Language models compare and draft with citations. Agents coordinate those services inside the workflow state and authority perimeter.

5.4 Data and evidence layer

The data layer includes master entities, documents, structured facts, market data, financial data, portfolio data, communication records and reference taxonomies. Each object carries ownership, classification, lineage, quality and retention attributes.

The evidence graph links source, fact, assumption, calculation, model, conclusion, term, decision and monitoring event. The same graph supports backward traceability from an external statement to its basis and forward traceability from a changed source to affected decisions.

5.5 Integration and event layer

APIs and governed event streams connect authoritative systems. The joint UAE supervisory guidelines address APIs, cloud computing, biometrics, big-data analytics, AI and distributed-ledger technology. They call for risk

management, security, data management, due diligence, outsourcing assessment and monitoring within their stated institutional scope [8].

An event contract specifies producer, consumer, object identifier, schema, timestamp, classification, version and delivery expectation. Consumers process events idempotently and record failures. Sensitive payloads are minimised; a permitted reference can replace data duplication.

5.6 Security and observability planes

Security and observability cross every layer. Identity, device posture, authorisation, encryption, secrets, data-loss controls and network policy apply to users, services and agents. Logs capture access, retrieval, prompts or task specifications, tool calls, model versions, validations, approvals, releases, latency, cost and incidents.

NIST SP 800-207 frames zero trust around resource protection and least-privilege access decisions for users, applications, services and devices [14]. This paper applies that principle at each request. A successful login does not create unrestricted investment-data access.

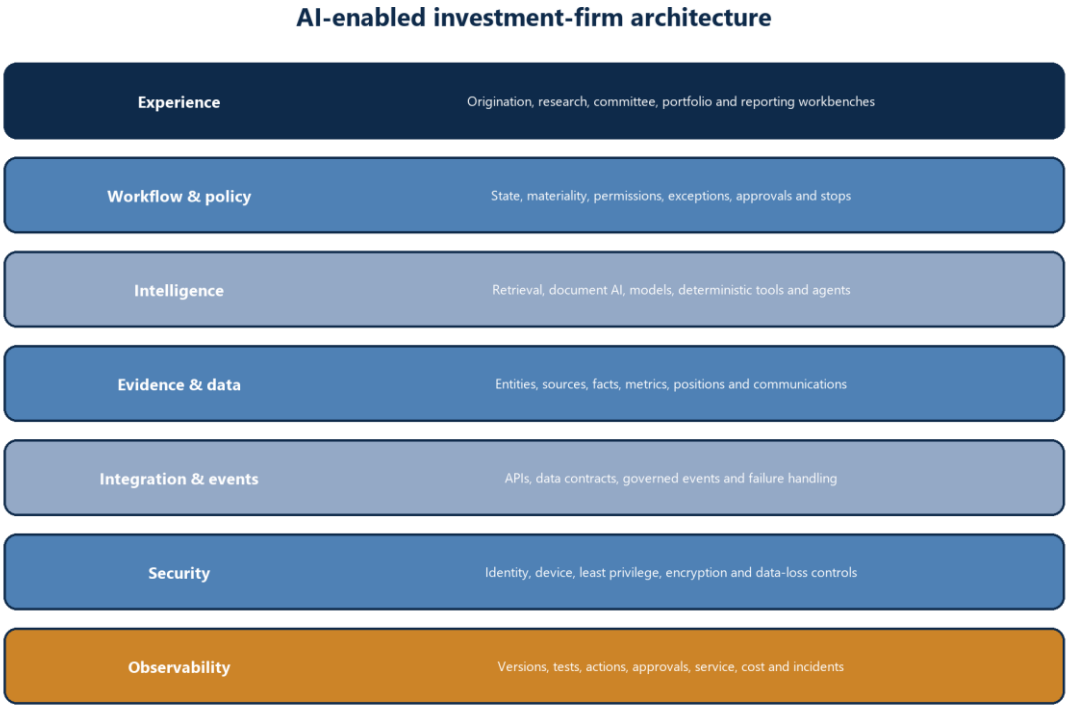

Figure 2. Layered architecture for an AI-enabled investment firm across experience, workflow, intelligence, evidence, integration, security and observability

Source: Matchpoint architecture; governance and technology references [4,8,11-15].

6. DATA AND EVIDENCE FOUNDATION

6.1 Entity graph

Investment relationships span people, organisations, funds, vehicles, assets, instruments, mandates, accounts and service providers. Names vary across documents and systems. An entity graph records stable internal identifiers, verified legal identifiers where available, roles, ownership, relationships, jurisdiction and effective dates.

Entity resolution should create proposed matches with confidence and evidence. A person approves ambiguous matches. The graph should preserve former names and historical relationships rather than overwrite them.

The entity graph supports conflict checks, exposure aggregation, related-party analysis, recipient validation and relationship intelligence. It also limits data mixing between clients, funds and mandates.

6.2 Evidence ledger

Every material fact should link to a source record containing file or system identity, owner, date, version, hash or system revision, reporting period, permission and supersession status. Proposed extracted facts retain exact source locations.

Evidence status	Meaning	Permitted use
Proposed	Extracted or generated for review	Internal preparation only
Verified	Checked against an approved source or reconciliation	May support analysis within scope
Disputed	Conflicting evidence remains	Requires visible exception
Superseded	Newer approved evidence exists	Historical trace only
Restricted	Use limited by permission, agreement or law	Permitted purpose and audience only
Expired	Currentness threshold exceeded	Refresh or documented exception required

The status remains attached when a fact appears in a model input, paragraph, chart or communication.

6.3 Taxonomy and semantic layer

A common taxonomy defines sectors, strategies, instruments, stages, decision types, risks, requirements and services. A semantic layer defines metrics and dimensions used across analysis and reporting.

For example, assets under management, committed capital, invested capital, net asset value and gross asset value are distinct. Revenue may be reported, run-rate, contracted or collected. Exposure may be gross, net, look-through or risk-weighted. The definition includes formula, unit, time basis, perimeter, owner and approved use.

Language models can map source terms to proposed canonical terms. Deterministic validation and owner approval establish the released mapping.

6.4 Data quality contracts

A data contract specifies schema, definition, source, frequency, quality checks, cut-off, owner and failure action. Quality dimensions include completeness, validity, reconciliation, timeliness, uniqueness and consistency.

A failed check creates an exception and identifies downstream consumers. A portfolio dashboard should show the freshness and coverage of its data, not only the calculated value.

6.5 Privacy and confidentiality

Investment firms may process personal data, commercially sensitive information, material non-public information, client records and restricted fund data. The architecture classifies information at source and enforces purpose, jurisdiction, retention and sharing rules.

DIFC's data-protection framework establishes obligations concerning collection, handling, use, individual rights and data transfers for entities within its scope [16]. The CBUAE's 2026 guidance addresses transparency, accountability, explainability, data privacy and human oversight for licensed financial institutions where AI may bear on consumers [7]. Each firm's legal analysis should identify the applicable regime and contract.

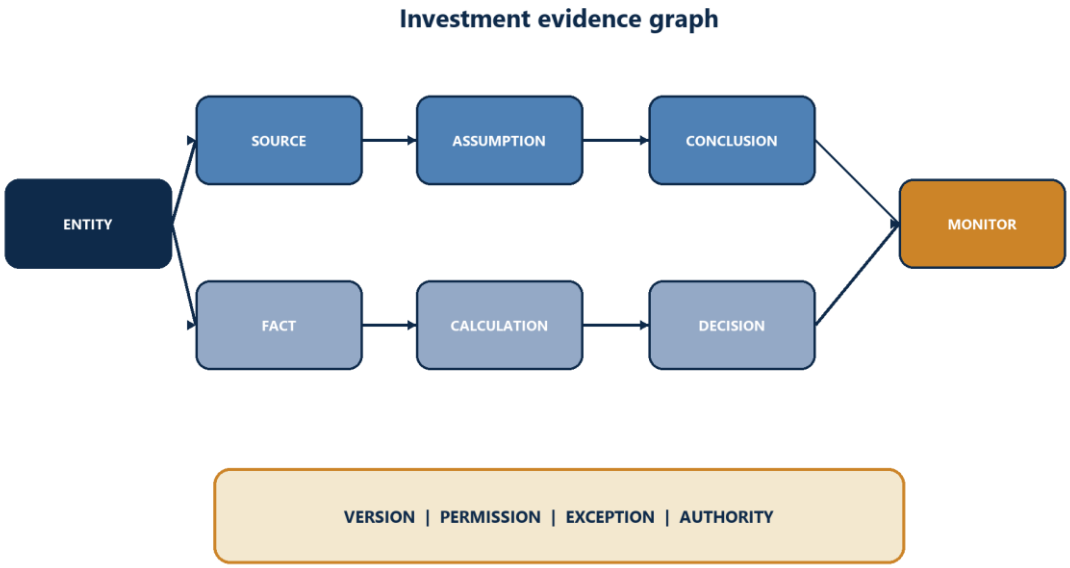

Figure 3. Investment evidence graph linking entity, source, fact, assumption, calculation, conclusion, decision and monitoring event

Source: Matchpoint framework.

7. MODELS, LANGUAGE SYSTEMS AND AGENTS

7.1 Use-case inventory

Every AI or model use case should have a registered purpose, owner, users, affected parties, data, output, materiality, provider, version, dependencies, intended use, prohibited use, tests, fallback and decision authority.

The inventory should include externally supplied scores, embedded vendor features and end-user-built tools. The Bank of England and FCA survey's third-party and understanding results support the need to inventory provider components and knowledge gaps [1].

7.2 Risk tiering

Risk tiering considers consequence, autonomy, data sensitivity, reversibility, scale, external exposure and detectability.

Tier	Example	Required gate
1; assistive	Internal search over approved policy	Source and permission tests
2; analytical	Proposed extraction or research classification	Field evaluation and reviewer
3; decision support	Investment scoring or scenario recommendation	Independent validation and authority boundary
4; material action	External communication, order, term or commitment	Named approval and release record

An output can move between tiers through its use. A summary used for orientation may be Tier 1. The same text quoted in an investor report becomes a material external statement and receives a higher release gate.

7.3 Retrieval and grounding

Permission-aware retrieval should filter sources before model access. The system records query, source set, versions, retrieved passages and citations. Retrieval quality is tested through labelled questions, expected sources, recall, precision and abstention.

A citation establishes source location. It does not establish that the claim is complete, current or correctly interpreted. Claim review tests support, contradiction, omission, scope and materiality.

7.4 Deterministic tools

Financial arithmetic, fee calculations, exposure aggregation, reconciliations and rule tests should use deterministic services where formulas can be specified. The service records code version, inputs, unit, currency, period, output and validation.

A language model can prepare inputs, explain a result and identify exceptions. The approved calculation service remains authoritative.

7.5 Statistical models

Statistical models require development data, outcome definitions, sampling, validation, calibration, stability, explainability, limitations and monitoring appropriate to intended use. Performance should be segmented where population differences matter.

A model release records the approved version and use boundary. A material data, feature, provider or objective change triggers reassessment. Human override records reason, authority and later outcome review.

7.6 Language models

Language systems are evaluated across tasks rather than through one global score. Tests cover source fidelity, unsupported material claims, completeness, instruction following, confidentiality, prompt injection, tool use, format, latency and cost.

NIST's AI RMF organises voluntary risk work around Govern, Map, Measure and Manage [11]. The Generative AI Profile adds risks and actions specific to generative systems, including confabulation, information integrity, privacy, security, human-AI configuration and third-party risk [12]. The paper maps these outcomes to the firm-level inventory, evaluation, release and incident processes.

7.7 Workflow agents

Provider guidance commonly defines an agent through a model, tools and instructions, with guardrails and evaluation around the workflow [19]. An investment-firm implementation should add explicit state, identity, source perimeter, budget, stop conditions and named authority.

An agent may retrieve approved evidence, call a calculation service, prepare a draft and open an exception. External transmission, model changes, transactions, binding terms and investment decisions remain behind authority gates appropriate to the firm.

7.8 Evaluation stack

Layer	Failure mode	Test	Release evidence
Evidence	Wrong source, version or permission	Manifest and access tests	Approved source set
Extraction	Wrong field, unit, sign or period	Exact-match and reconciliation tests	Verified structured facts
Calculation	Mapping or formula error	Independent recomputation	Calculation pass
Model	Misuse, bias or deterioration	Validation and monitoring	Approved intended use
Language	Unsupported or incomplete claim	Citation, claim and omission review	Cited accepted draft
Workflow	Wrong state, tool or recipient	State, stop and authority tests	Approved release record

Figure 4. Evaluation and authority stack from evidence through component tests, workflow controls and human release
Source: Matchpoint framework; risk and conduct references [2-5,11,12].

8. INVESTMENT-LIFECYCLE APPLICATIONS

8.1 Origination and relationship intelligence

The platform can unify approved relationship, mandate, market and interaction evidence. It can propose opportunity classifications, identify relationship paths, prepare meeting briefs and track next actions.

Controls include lawful and permitted data use, source currentness, conflict checks, duplicate resolution, communication approval and opt-out handling. A generated relationship claim should cite its basis.

8.2 Research and thematic intelligence

The research workbench registers sources, extracts proposed facts, groups evidence by thesis, identifies disagreement and drafts cited notes. A thesis object records supporting and contrary evidence, assumptions, time horizon, catalysts, risks and falsification conditions.

The system should preserve publication date, event date and research cut-off. Market data and company data retain provider rights. A language model should abstain when the approved source set cannot support the question.

8.3 Screening and qualification

Deterministic rules test mandate fit, geography, sector, ticket, stage, return, concentration and exclusions. Models may rank opportunities inside an approved population. The output shows reasons, missing data and exceptions.

The decision remains a named qualification event. A high score does not bypass a conflict, legal restriction or capacity constraint.

8.4 Diligence and underwriting

Document services inventory evidence, compare versions, extract proposed facts and track requests. Calculation services build approved ratios and scenarios. Language models create cited issue summaries. Specialist reviewers own conclusions.

The final decision object connects commercial, financial, technology, legal, operational and compliance findings. A material issue has severity, evidence, owner, mitigation, residual risk and decision effect.

8.5 Investment committee

The committee workbench presents the proposed decision, thesis, economics, downside, sensitivities, exceptions, unresolved matters, challenge and conditions. It highlights changes since the previous version.

The approval record captures forum, attendance, declared conflicts, recommendation, vote or authority path, conditions, dissent and monitoring consequences. AI can prepare and compare; the committee owns the decision.

8.6 Portfolio monitoring and valuation

Approved underwriting assumptions, terms and triggers flow into the portfolio record. Monitoring ingests financial, operational, covenant, market and qualitative evidence. It distinguishes reported values, calculated values, model outputs and manager judgement.

Events identify changes from plan, threshold breaches and missing reports. The system routes them to the named owner and authority. Valuation conclusions retain source, method, assumptions, challenge and approval.

8.7 Investor relations and fundraising

The platform connects approved fund data, performance, pipeline, due-diligence responses, relationship state and communication permissions. It can prepare tailored drafts from an approved disclosure library and current data cut-off.

Claims about track record, pipeline, technology, team, ESG or operations should retain a basis and approval record. The SEC AI-washing orders demonstrate the specific enforcement risk of unsupported technology representations for registered advisers [6].

8.8 Reporting and client service

Reporting uses governed metrics, approved commentary and versioned templates. Every number links to its definition and source. Every material narrative claim links to evidence and review.

ESMA's 2024 statement emphasises that management bodies remain responsible for firms' decisions when AI supports retail investment services within MiFID II scope [5]. The paper applies the same accountability design principle to external investment and advisory communications while retaining the statement's jurisdictional boundary.

9. SECURITY, PRIVACY AND RECORDS

9.1 Identity-centred access

Access decisions consider user or service identity, role, purpose, device, data classification, jurisdiction, time and workflow state. Permissions apply to retrieval, calculation, drafting, writing and release separately.

Service identities and agent identities receive least privilege. Short-lived credentials and approved secrets management reduce persistent access. Privileged actions require stronger authentication and logging.

9.2 Data-loss controls

Controls can restrict uploads to unapproved models, copying of sensitive data, external sharing, local downloads and use of personal accounts. Documents are classified at intake. Sensitive fields can be masked or tokenised where the task allows.

Provider retention, training use, sub-processors, hosting, encryption, deletion and incident terms require explicit review. A contract statement should be tested against the actual service configuration.

9.3 Secure development and change

AI-enabled workflows include prompts, tools, retrieval configurations, schemas, code, models and policies. All are change-controlled assets. Development and production environments are separated. Tests cover authentication, authorisation, dependency risk, prompt injection, data exfiltration, malicious files, tool misuse and rollback.

NIST SP 800-218A extends secure software-development considerations to generative AI and dual-use foundation models [15]. The paper applies its lifecycle orientation to the firm's AI platform and vendor integration.

9.4 Records and communications

The record policy maps activity, jurisdiction and requirement to retained objects. It may include source evidence, calculations, approvals, communications, advertisements, client records, model versions, test results and incidents.

The SEC's electronic recordkeeping rules for funds and advisers permit electronic retention under conditions including safeguarding records from loss, alteration or destruction, limiting access and maintaining complete, true and legible copies within their scope [18]. Firms should obtain current legal advice for their activities and jurisdictions.

Generated communication channels should route through approved archives. A draft in a model interface should not become an unrecorded client or investor communication.

9.5 Cyber risk and incident management

NIST CSF 2.0 organises cybersecurity outcomes around Govern, Identify, Protect, Detect, Respond and Recover [13]. DFSA's cyber-supervision summary expects firms within its scope to establish a cyber strategy and framework tailored to nature, scale and complexity [17].

The incident record captures affected data, workflow, model, provider, output, users, recipients, containment, recovery, notification, root cause and corrective action. AI incidents include unsupported material claims, confidential-data exposure, unauthorised tool action, corrupted retrieval, calculation mismatch, harmful bias, provider outage and monitoring failure.

Figure 5. Investment-firm risk-control map across data, models, agents, cyber security, third parties, records and authority
Source: Matchpoint synthesis based on [4,7-18].

10. GOVERNANCE AND DECISION RIGHTS

10.1 Board and senior-management mandate

The governing body or senior accountable forum approves the AI and technology mandate, risk tolerance, materiality criteria, prohibited uses, funding, accountability and reporting. The mandate identifies the business outcomes sought and the decisions that remain reserved.

The FSB's June 2026 consultation proposes 12 sound practices across organisation-wide AI governance and the AI lifecycle for financial institutions [4]. The document is a consultation. Its themes support a governance design covering strategy, accountability, skills, data, models, third parties, monitoring and incidents.

10.2 Authority matrix

Activity	System or AI	Business owner	Independent control	Decision authority
Register source	Prepare record	Verify owner and scope	Sample or challenge	Informed
Propose fact	Extract and cite	Verify	Test quality	Informed
Run approved calculation	Execute service	Review inputs	Recompute sample	Use within scope
Produce model output	Execute approved model	Interpret	Validate and monitor	Consider within boundary
Draft analysis	Prepare cited draft	Own conclusion	Challenge material claims	Review or approve
Propose external message	Prepare from approved evidence	Confirm recipient and purpose	Compliance review where required	Named approver releases
Propose transaction or term	Prepare analysis	Recommend	Risk or legal challenge	Committee or delegated authority
Change model or policy	Prepare change package	Sponsor	Independent validation	Change authority approves

The exact allocation depends on firm governance. The invariant is that system capability does not change delegated authority.

10.3 Change approval

A change package records objective, affected use cases, data, model or tool version, tests, limitations, migration, fallback, communications and owner. Material changes receive independent review and approval before production.

Emergency changes have a narrow duration, named authority, enhanced logging and post-event review. A provider's silent model update should be treated as a dependency change when it may affect output.

10.4 Monitoring and reporting

Management information should cover:

- registered and unregistered use cases;
- risk-tier distribution;
- data and source-quality failures;
- evaluation results and drift;
- accepted-output and reviewer effort;
- overrides, exceptions and abstentions;
- authority violations and blocked actions;
- provider concentration and outages;

- incidents and remediation;
- cost by capability and workflow; and
- realised outcome measures where attribution is defensible.

Counts of generated documents or model calls do not establish business value. Accepted work, decision quality, service resilience, cost and risk evidence provide a stronger basis.

11. THIRD-PARTY AND OPERATIONAL RESILIENCE

11.1 Provider inventory

The provider inventory covers cloud, data, model, embedding, document-processing, workflow, portfolio, CRM, communication and security services. It records legal entity, service, region, data, sub-processors, dependency, criticality, concentration, contract, service levels, audit rights, change notice, exit and fallback.

The Bank of England and FCA survey reported that the top three named providers accounted for 73% of cloud providers, 44% of model providers and 33% of data providers among respondents [1]. These respondent-reported concentrations support explicit dependency mapping.

11.2 Due diligence and contracting

The joint UAE enabling-technology guidelines call for materiality and outsourcing risk assessment, provider due diligence, security controls, data-location consideration, monitoring and exit planning for institutions within scope [8]. The architecture translates these themes into a provider evidence pack.

The pack includes security and privacy evidence, financial and operational condition, sub-contracting, data rights, model update policy, incident obligations, business continuity, recovery, testing rights and transition support.

11.3 Important services and impact tolerances

The FCA's operational-resilience framework requires in-scope firms to identify important business services, set impact tolerances and map and test their ability to remain within those tolerances. The transition deadline for in-scope firms was 31 March 2025 [10]. The paper uses service mapping and testing as a design reference; it does not represent every investment firm as within scope.

An investment firm can identify services such as order execution, valuation, capital calls, investor reporting, client communication or covenant escalation. Each service maps people, process, data, systems, providers and facilities. Severe but plausible tests examine loss of a model provider, data feed, cloud region, identity service, portfolio platform or key person.

11.4 DORA context

The EU Digital Operational Resilience Act applies within its scope to financial entities including investment firms, managers of alternative investment funds and management companies, subject to stated exclusions. It requires an internal governance and control framework for ICT risk and assigns management-body responsibility [9]. Applicability requires legal analysis.

The architecture supports the relevant design themes through asset and provider inventories, incident records, testing, change control, service mapping and exit plans.

11.5 Exit and portability

Exit planning identifies data export, schema, model replacement, workflow fallback, access revocation, retention, deletion verification, intellectual property and transition time. A critical capability should have a tested degraded mode.

Portability does not require every component to be interchangeable immediately. It requires visibility into dependency and a credible transition path proportional to criticality.

12. PRODUCTIVITY AND ECONOMIC CASE

12.1 Measurement unit

The measurement unit is an accepted work packet or completed decision object. A packet counts when required evidence, calculation, exceptions, review and authority records pass its release gate.

The baseline should observe:

- attempted and accepted packets;
- analyst and reviewer active time;
- total elapsed time;
- source coverage and citation validity;
- calculation and reconciliation defects;
- first-pass acceptance and rework;
- exceptions and abstentions;
- late or incorrect releases;
- provider, data and control cost; and
- incidents and remediation.

12.2 Value equation

Let:

- (N) be attempted packets per year;
- (a) be the accepted-output rate;
- (H_b) be baseline active hours per accepted packet;
- (H_p) be post-pilot active hours per accepted packet;
- (c_h) be loaded hourly cost;
- (C_t) be annual technology and data cost;
- (C_v) be annual validation, security and control cost;
- (C_r) be observed rework cost; and
- (C_i) be observed incident and remediation cost.

The illustrative operating-value structure is:

Annual operating value = N × a × (H_b - H_p) × c_h - C_t - C_v - C_r - C_i

Revenue or investment-performance value should remain zero unless the firm has a defined causal mechanism, an observation period, a valid comparison and approved attribution.

12.3 Illustrative scenario

The following scenario is unverified and uses illustrative management assumptions solely to demonstrate the calculation structure.

Input	Bounded pilot	Multi-workflow platform	Scaled firm	Status
Attempted packets per year	300	1,200	3,000	Illustrative management assumption
Accepted-output rate	65%	80%	88%	Illustrative management assumption

Input	Bounded pilot	Multi-workflow platform	Scaled firm	Status
Baseline active hours per accepted packet	8	10	12	Illustrative management assumption
Post-pilot active hours per accepted packet	6.5	7.5	8.5	Illustrative management assumption
Loaded hourly cost	USD 120	USD 160	USD 200	Illustrative management assumption
Technology and data cost	USD 100,000	USD 350,000	USD 900,000	Illustrative management assumption
Validation and control cost	USD 80,000	USD 250,000	USD 600,000	Illustrative management assumption
Approved revenue or investment value	USD 0	USD 0	USD 0	No observed basis assumed

This table does not establish a business case. It omits any forecast of better investment selection, fundraising conversion, retention, fees or portfolio performance. Those outcomes require a longer observation period, comparable populations and a defensible attribution method.

12.4 Productivity tree

Productivity has four branches.

- **Accepted capacity:** analyst hours, reviewer hours and throughput.
- **Evidence quality:** coverage, citation validity, calculation fidelity and completeness.
- **Decision service:** elapsed time, exception closure, approval readiness and stakeholder experience.
- **Cost and risk:** technology, data, controls, rework, incidents and provider concentration.

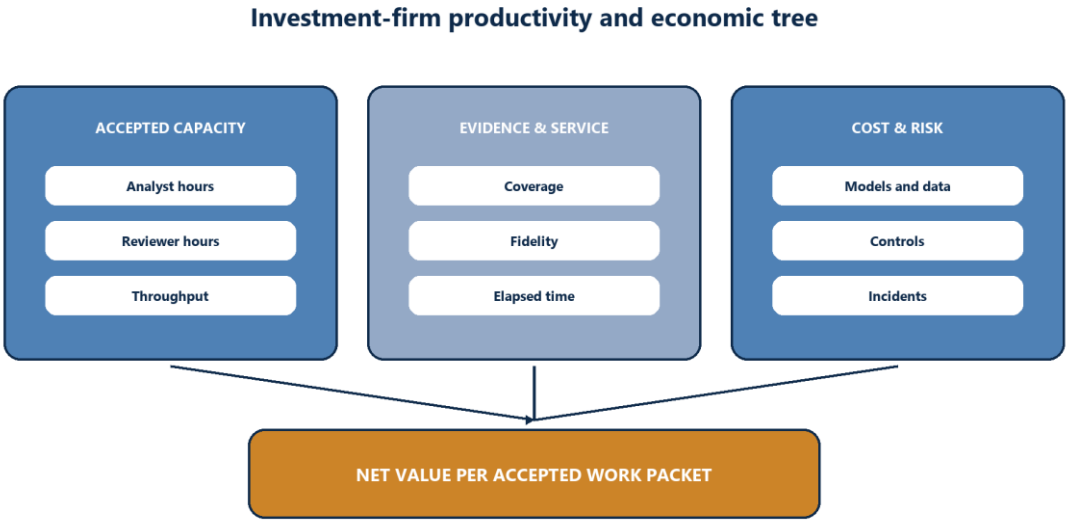

Figure 6. Productivity and economic tree for accepted capacity, evidence quality, decision service, cost and risk
Source: Matchpoint scenario analysis. Operating inputs require observed baseline and pilot evidence.

13. IMPLEMENTATION ROADMAP

13.1 Stage 1; mandate and inventory

Define business outcomes, scope, services, decisions, authority, jurisdictions, data classes, providers and current systems. Inventory existing AI and end-user tools. Select one bounded workflow with sufficient volume and historical examples.

Exit gate: approved charter, owner, risk tier, source perimeter, baseline and prohibited actions.

13.2 Stage 2; identity and evidence foundation

Create stable entity identifiers, source registration, permission-aware retrieval, taxonomy, metric definitions and evidence status. Establish deterministic calculation and reconciliation services for the pilot.

Exit gate: source-to-output traceability and permission tests pass on the labelled evaluation set.

13.3 Stage 3; assistive pilot

Deploy read-only search, classification, proposed extraction and cited drafting. Reviewers remain blind to model origin where practical during comparative evaluation. Record omissions, unsupported claims, acceptance and effort.

Exit gate: component thresholds pass for the intended document population; excluded cases remain excluded.

13.4 Stage 4; workflow shadowing

Run the full work packet in parallel with the current process. Test workflow state, exceptions, deterministic tools, review, approval and records. Compare accepted output and total effort.

Exit gate: complete-packet quality is non-inferior under the approved test and authority remains unchanged.

13.5 Stage 5; controlled production

Release selected internal outputs behind explicit approvals. Monitor data failures, unsupported claims, tool errors, overrides, incidents, cost and elapsed time. Maintain a tested fallback.

Exit gate: the control forum approves continued use from observed production evidence.

13.6 Stage 6; platform integration

Connect the approved workflow to systems of record, portfolio data, communications and management reporting. Establish service levels, recovery tests and provider exit evidence.

Exit gate: end-to-end service, record, resilience and data-transfer tests pass.

13.7 Stage 7; expansion

Add new decisions, teams, strategies or jurisdictions through a new scope and validation package. Reuse common services and retain separate use-case evidence.

Exit gate: the change authority approves scope, test evidence, controls and operating ownership.

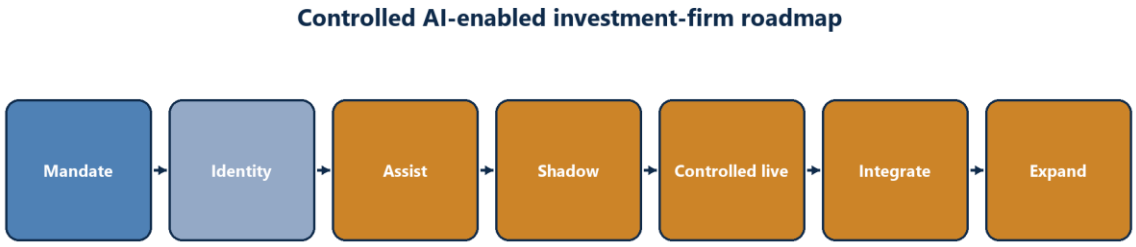

Figure 7. Controlled adoption roadmap from mandate and evidence foundation through assistive use, shadowing, production, platform integration and expansion

Source: Matchpoint framework.

14. APPLICATION TO FAMILY OFFICES AND FUND MANAGERS

14.1 Family-office investment team

An A2 family-office team may oversee public markets, funds, direct deals, private credit, real estate, liquidity and operating businesses. The architecture can provide a consolidated entity and evidence graph while preserving mandate, vehicle, family-member and adviser permissions.

Priority decision objects may include:

- manager due diligence and annual review;
- direct-deal screening and committee packs;
- portfolio cash-flow and commitment pacing;
- exposure and concentration analysis;
- document and tax-record organisation;
- family or investment-committee reporting; and
- monitoring of covenants, valuations and material events.

The family office should define which information may be shared across family members, entities, employees and external advisers. The human CIO or committee retains allocation and investment authority.

14.2 Fund manager or general partner

A B2 manager may prioritise fundraising, investor diligence, pipeline, underwriting, portfolio operations, valuation, reporting and compliance. The architecture can reuse approved fund data and disclosure evidence across investor workflows while separating prospective-investor permissions and communication state.

Priority decision objects may include:

- target-LP qualification;
- meeting and diligence preparation;
- DDQ and data-room response;
- investment research and committee packs;
- portfolio review and valuation;
- quarterly reporting and letters;
- material-event communication; and
- marketing-claim substantiation.

The manager should connect technology representations, track-record metrics and performance commentary to approved evidence. Investor communication remains behind the applicable review and record process.

14.3 Shared architecture opportunity

Family offices and managers can align around permissioned decision objects. A manager may publish an approved diligence packet with definitions, evidence dates, methodology, exceptions and updates. An allocator can ingest that packet without receiving unrestricted internal data. Stable identifiers and data contracts reduce repeated re-keying while preserving each party's authority.

15. LIMITATIONS AND RESEARCH AGENDA

First, the cited adoption surveys aggregate different financial-services sectors. They do not establish use, capability or productivity for a specific family office, fund manager or advisory firm [1,3].

Second, consultation reports from IOSCO and the FSB may change before final publication and do not create binding standards [3,4].

Third, legal and regulatory obligations vary by jurisdiction, entity, client and activity. The control frameworks cited here cannot replace current legal analysis.

Fourth, provider models, product configurations and data practices change. A test result applies to the evaluated version, data, tools and workflow.

Fifth, a small firm may combine roles. The requirement for independent challenge should be implemented proportionately and documented.

Sixth, the numerical scenario is unverified and uses illustrative management assumptions. It supports calculation design only.

Future research should build permissioned benchmark packets for origination, research, diligence, committee preparation, monitoring and reporting. Each benchmark should contain realistic sources, changed versions, permissions, calculations, exceptions, approved outputs and failure cases. Evaluation should measure component fidelity, complete-packet acceptance, reviewer time, authority compliance, cost and incidents. Longitudinal study is required before linking the architecture to fundraising, investment or portfolio outcomes.

16. CONCLUSION

An AI-enabled investment firm requires a governed decision architecture. Stable entities connect relationships, funds, deals, assets and communications. Registered evidence connects each material fact to its source and status. Deterministic services perform reproducible calculations. Statistical models operate within approved use. Language models retrieve and draft with citations. Workflow agents coordinate tasks inside explicit state, tools, budgets and stop conditions. Named people retain material decision and release authority.

The architecture joins commercial, investment, control and management systems through decision objects and governed events. Systems of record remain authoritative for their objects. Security and observability cross every layer. Provider dependencies, service resilience, data rights, records and exit receive design treatment from the beginning.

Implementation begins with a bounded mandate and an observed baseline. It establishes identity and evidence, tests components, shadows complete workflows, releases controlled internal outputs and expands after observed evidence. Value is measured through accepted work packets, evidence quality, reviewer effort, service, cost and incidents.

For family offices and fund managers, the practical objective is a more complete and timely investment process whose evidence, judgement and authority remain visible from origination to portfolio reporting.

DECLARATIONS

Author. Chennakeshav (CK) Adya.

Affiliation. Independent Researcher.

Conflict of interest. The author is the Founder and Managing Partner of Matchpoint Partners. The paper describes a Matchpoint framework and does not report a client engagement or product performance.

Funding. No external funding was received for this paper.

Data availability. The paper uses public sources listed below. The numerical economic scenario is generated for framework demonstration and is not observed firm data.

Use of AI. AI tools assisted research organisation, drafting, document production and quality checks under author direction. The author retains responsibility for the paper.

REFERENCES

[1] Bank of England and Financial Conduct Authority. (2024). *Artificial intelligence in UK financial services: 2024*. https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024

[2] International Organization of Securities Commissions. (2021). *The use of artificial intelligence and machine learning by market intermediaries and asset managers: Final report*. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf

[3] International Organization of Securities Commissions. (2025). *Artificial intelligence in capital markets: Use cases, risks, and challenges: Consultation report*. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf

[4] Financial Stability Board. (2026). *Sound practices for financial institutions' responsible AI adoption: Consultation report*. https://www.fsb.org/uploads/P100626.pdf

[5] European Securities and Markets Authority. (2024). *Public statement on the use of artificial intelligence in the provision of retail investment services*. https://www.esma.europa.eu/sites/default/files/2024-05/ESMA35-335435667-5924__Public_Statement_on_AI_and_investment_services.pdf

[6] U.S. Securities and Exchange Commission. (2024). *SEC charges two investment advisers with making false and misleading statements about their use of artificial intelligence*. https://www.sec.gov/newsroom/press-releases/2024-36

[7] Central Bank of the UAE. (2026). *Guidance note on the consumer protection and responsible adoption and use of artificial intelligence and machine learning by licensed financial institutions in the U.A.E.* https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence

[8] Central Bank of the UAE, Securities and Commodities Authority, Dubai Financial Services Authority and Financial Services Regulatory Authority. (2021). *Guidelines for financial institutions adopting enabling technologies*. https://www.adgm.com/documents/legal-framework/guidance-and-policy/fsra/guidelines-for-financial-institutions-adopting-enabling-technologies-20211107.pdf

[9] European Parliament and Council. (2022). *Regulation (EU) 2022/2554 on digital operational resilience for the financial sector*. https://eur-lex.europa.eu/eli/reg/2022/2554/oj

[10] Financial Conduct Authority. (2026). *Operational resilience*. https://www.fca.org.uk/firms/operational-resilience

[11] National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

[12] National Institute of Standards and Technology. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1. https://doi.org/10.6028/NIST.AI.600-1

[13] National Institute of Standards and Technology. (2024). *The NIST Cybersecurity Framework (CSF) 2.0*. NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

[14] Rose, S., Borchert, O., Mitchell, S. and Connelly, S. (2020). *Zero Trust Architecture*. NIST SP 800-207. https://doi.org/10.6028/NIST.SP.800-207

[15] National Institute of Standards and Technology. (2024). *Secure Software Development Practices for Generative AI and Dual-Use Foundation Models: An SSDF Community Profile*. NIST SP 800-218A. https://doi.org/10.6028/NIST.SP.800-218A

[16] Dubai International Financial Centre Commissioner of Data Protection. (2020, amended). *Guide to Data Protection Law, DIFC Law No. 5 of 2020 and Data Protection Regulations*. https://www.difc.com/business/registrars-and-commissioners/commissioner-of-data-protection

[17] Dubai Financial Services Authority. (2026). *Cyber risk supervision: Summary*. https://www.dfsa.ae/what-we-do/supervision/cyber-risk-supervision/summary

[18] U.S. Securities and Exchange Commission. (2001). *Electronic recordkeeping by investment companies and investment advisers*. https://www.sec.gov/rules-regulations/2001/05/electronic-recordkeeping-investment-companies-investment-advisers-correction

[19] OpenAI. (2025). *A practical guide to building agents*. https://cdn.openai.com/business-guides-and-resources/a-practical-guide-to-building-agents.pdf

APPENDIX A. DECISION OBJECT SCHEMA

Field	Example requirement
Decision ID	Stable identifier and decision type
Objective	Defined question and required outcome
Entity and instrument	Verified internal and legal identifiers
Mandate	Strategy, vehicle, client and policy perimeter
Source set	Registered evidence, versions, permissions and cut-off
Facts	Proposed, verified, disputed, superseded or restricted
Assumptions	Value, range, basis, owner and scenario
Calculations	Formula or code version, inputs, output and validation
Models	Model, version, intended use, result and limitation
Exceptions	Policy or evidence issue, severity, owner and resolution
Recommendation	Proposed action, rationale and conditions
Challenge	Independent questions, response and residual issue
Decision	Forum, authority, date, outcome, dissent and conditions
Monitoring	Metric, threshold, frequency, source and escalation

APPENDIX B. AI USE-CASE APPROVAL MATRIX

Output or action	Preparer	Validator	Approver	Release constraint
Source manifest	System or analyst	Business reviewer	Workstream lead	Approved source perimeter
Extracted fact	Document service	Analyst and reconciliation	Workstream lead	Verification status visible
Financial calculation	Deterministic service	Independent recomputation	Named analyst	Approved definition and inputs
Statistical score	Approved model	Model owner or validator	Risk authority	Intended use only
Research note	Analyst with AI support	Senior reviewer	Research authority	Cited claims and current cut-off
Investment memorandum	Deal team	Independent challenge	Committee	Sources, exceptions and conditions complete
Investor response	IR team with AI support	Data owner and compliance	Named approver	Approved recipient and disclosure set
Portfolio alert	Monitoring service	Portfolio owner	Escalation authority	Threshold and source recorded
Model or tool change	Product team	Security, privacy and validation	Change authority	Tests, fallback and version record

APPENDIX C. PILOT SCORECARD

Dimension	Measure	Definition
Evidence	Source coverage	Required sources registered and current

Dimension	Measure	Definition
Evidence	Citation validity	Material claims supported by exact source locations
Extraction	Field exact match	Value, unit, period and perimeter all correct
Calculation	Reconciliation pass	Statements, schedules and formulas satisfy approved checks
Model	Calibration or task metric	Performance meets approved intended-use threshold
Model	Stability	Performance and data drift remain within limits
Language	Unsupported-claim rate	Material claims without adequate evidence
Workflow	Tool precision	Permitted tool calls with correct parameters
Workflow	Stop accuracy	Required abstentions and escalations occur
Authority	Authority violation	Material or external action without approval
Quality	First-pass acceptance	Packets accepted without material correction
Productivity	Reviewer time	Median active reviewer hours per accepted packet
Productivity	Accepted throughput	Accepted packets per defined period
Service	Elapsed time	Time from accepted intake to released decision object
Risk	Incident rate	Classified incidents per production packet
Cost	Total operating cost	Models, data, tools, people, validation and controls

APPENDIX D. SERVICE AND PROVIDER RECORD

Minimum fields include provider legal entity, service, owner, business capability, data classes, locations, sub-processors, identity method, permissions, model or product version, criticality, concentration, service level, recovery objective, audit evidence, change notice, incident contact, contract dates, termination rights, data export, deletion verification, replacement option and last test.

APPENDIX E. GLOSSARY

Accepted work packet. A bounded task whose evidence, output, checks, exceptions, review and authority records satisfy the defined release gate.

Decision object. The governed record connecting an investment or advisory question to evidence, assumptions, calculations, models, challenge, authority and monitoring consequences.

Deterministic service. A rule or calculation that produces the same output from the same approved inputs and version.

Evidence graph. The linked record of entities, sources, facts, assumptions, calculations, models, conclusions, terms, decisions and monitoring events.

Intended use. The approved decision, population, data and operating context for a model or system.

System of record. The authoritative repository for a defined business object and its approved state.

Workflow agent. A system that coordinates retrieval, tools and state transitions within explicit permissions, budgets, stop conditions and authority.