MatchPoint Research Working Paper No. 2026/T15

AI Document Intelligence: Turning Data Rooms into Decisions

A controlled evidence architecture for family-office CIOs and GCC fund managers

Prepared by Matchpoint Partners
Author attribution pending CK approval
August 2026

ABSTRACT

Background. Private-market decisions depend on evidence distributed across agreements, questionnaires, financial statements, track-record workbooks, policies, filings and investor responses. **Objective.** This paper develops a controlled architecture for AI document intelligence in family-office diligence and GCC fundraising. **Approach.** The analysis reviews 40 official industry, legal, regulatory, standards and primary technical sources available through 1 August 2026. **Findings.** Source identity, version precedence, evidence coordinates, typed facts, deterministic calculations, conflict handling, abstention and named review should remain attached to every released answer. The accountable unit is a quality-adjusted accepted decision-evidence packet. **Implications.** A2 teams can organise manager diligence and committee evidence; B2 teams can improve data-room readiness and controlled investor responses. Attributed revenue, cost reduction and loss reduction remain USD 0 until approved observed attribution exists.

Highlights

Unit. A quality-adjusted accepted decision-evidence packet connects productivity to reproducible evidence.

Lineage. Every material fact retains source version, page or cell coordinates, context and extraction history.

Abstention. Missing, inaccessible, ambiguous, conflicting or out-of-authority questions are withheld and escalated.

Authority. Legal, accounting, regulatory, valuation, cybersecurity and investment decisions remain with qualified owners.

Economics. All scenarios are illustrative; commercial attribution remains zero without approved observed evidence.

JEL Classification: C55, C81, D81, G23, G24, L86, O33

Keywords: document intelligence, data rooms, private markets, family offices, fund managers, due diligence, retrieval, table extraction, evidence lineage, human review, AI governance

Research paper only. This document is not investment, legal, regulatory, privacy, cybersecurity, accounting, tax, audit, valuation or technology advice. All economic scenario inputs and pilot thresholds are unverified illustrative management assumptions.

1. INTRODUCTION

Private-market decisions are assembled from documents. A family-office investment committee may receive a private-placement memorandum, limited partnership agreement, side letters, due-diligence questionnaire, track-record workbook, valuation policy, audited accounts, portfolio schedule, regulatory filings, cybersecurity material and management responses. A fund manager may answer similar questions across several investors while controlling versions, permissions and representations. The volume is visible. The harder problem is preserving the relationship between every conclusion and the source that supports it.

AI document intelligence can classify files, extract fields, recover tables, retrieve passages, compare versions, calculate controlled metrics and draft evidence-linked answers. These capabilities can shorten mechanical search and transcription. They can also detach a number from its unit, period, vehicle, currency, gross-or-net basis, qualifying language or superseding version. A fluent answer can conceal an unanswerable question. A fast workflow can propagate a source conflict to an investment paper, investor response or marketing communication.

This paper develops an operating framework for **AI Document Intelligence: Turning Data Rooms into Decisions**. The primary ideal customer profile is A2, Family-Office CIOs & Heads of Alternatives: investment professionals allocating to funds and direct opportunities. The secondary profile is B2, GCC Fund Managers / GPs Raising Capital: emerging and established managers preparing and maintaining fund-raising evidence. These profiles and the title, hook, visual requirements and service mapping were verified from Matchpoint Partners' Topic Tracker and ICP Legend on 1 August 2026.

The accountable unit of productivity is a **quality-adjusted accepted decision-evidence packet**. The packet contains a controlled question, source-document identities, page or cell evidence spans, extracted facts, calculation lineage, conflicts, confidence, reviewer disposition and released answer. Pages processed, fields extracted, tokens generated and questions answered are operating statistics. Accepted packets connect effort to evidence that a named reviewer can reproduce and approve.

Seven propositions organise the analysis. First, source identity and version precedence come before extraction. Second, a retrieved passage or extracted field remains a candidate fact until its context and evidence span are preserved. Third, calculation and interpretation are separate authorities. Fourth, the system should abstain when evidence is absent, conflicting, inaccessible or outside the approved question. Fifth, permissions and confidentiality should follow the source into every index, cache, prompt, output and log. Sixth, validation should be field-, document-, question- and decision-specific. Seventh, productivity and commercial value should be measured through representative shadow pilots; attributed revenue, cost reduction and loss reduction remain USD 0 until approved observed attribution exists.

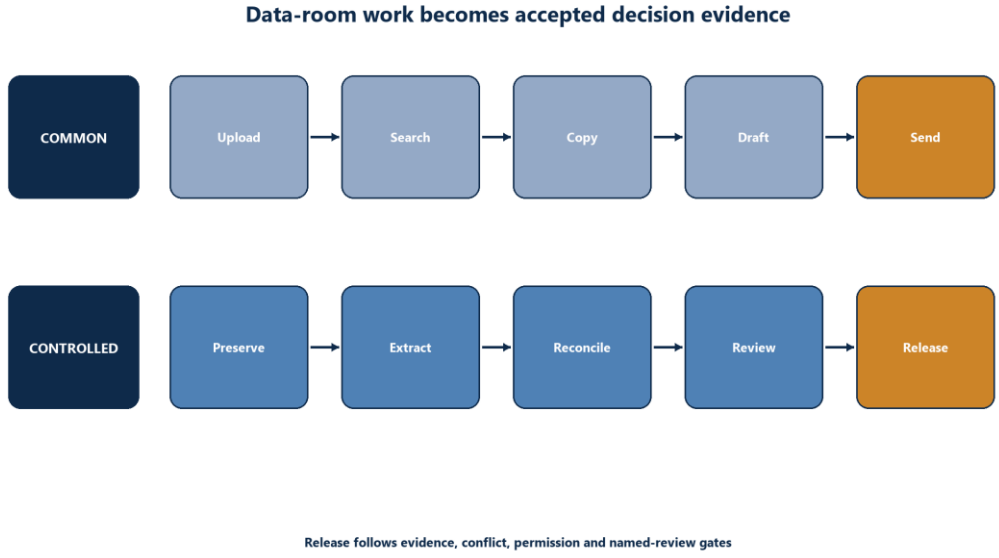

Figure 1. Before and after data-room workflow

Source: Matchpoint framework.

The evidence cut-off is 1 August 2026. The analysis uses official industry templates, primary laws, regulator material, official standards and primary technical research. It contains no Matchpoint client data room, fund document set, investor file, model output, observed workflow baseline or completed pilot. Worked volumes, time inputs, costs, acceptance rates and thresholds are explicitly unverified illustrative management assumptions. Legal, regulatory, privacy, cybersecurity, accounting, tax, audit, valuation and investment determinations remain with appropriately qualified owners.

The contribution is operational. Sections 2 and 3 define the evidence model and source hierarchy. Section 4 maps A2 and B2 use cases. Section 5 explains document-intelligence methods and failure modes. Section 6 defines retrieval, extraction, calculation and conflict controls. Section 7 specifies a controlled architecture. Section 8 sets out validation and human authority. Section 9 presents measurement and illustrative economics. Section 10 covers legal, regulatory, privacy and records boundaries. Sections 11 and 12 provide operating playbooks and a gated adoption roadmap. Section 13 states limitations and a research agenda. Appendices provide the evidence-packet schema, validation scorecard, procurement questions, release checklist and management questions.

2. SCOPE, DEFINITIONS AND EVIDENCE BOUNDARIES

2.1 What document intelligence means here

Document intelligence is the controlled conversion of documents into searchable, structured and reviewable evidence. It can combine optical character recognition, layout analysis, classification, table extraction, named-entity recognition, retrieval, question answering, comparison, calculation and generative drafting. The term covers scanned images, digitally generated PDFs, spreadsheets, presentations, word-processing files, emails and structured filings. Each format exposes different evidence and failure modes.

The paper distinguishes six objects. A source document is the preserved file received from an identified origin. A document version is one immutable state of that source. An evidence span is a page region, paragraph, table cell, worksheet cell or structured-data fact that supports a proposition. A candidate fact is a machine-extracted value linked to an evidence span. A calculation is a deterministic transformation of approved inputs. A released conclusion is a reviewer-approved statement used for a defined decision or communication.

Object	Minimum identity	Permitted initial use	Excluded status
Source document	File hash, origin, receipt time, access class and declared type	Preservation and catalogue	Current or authoritative merely because it exists
Document version	Source family, version ID, effective date and supersession status	Version comparison	Latest unless precedence is established
Evidence span	Document version, page/cell coordinates, text or image region and retrieval method	Candidate support	Standalone proof without context
Candidate fact	Value, unit, period, entity, basis, evidence span and extraction version	Review queue	Accepted fact
Controlled calculation	Approved input facts, formula, units, code version and result	Reproduction and review	Accounting, valuation or investment judgement
Released conclusion	Question, accepted evidence, reviewer, limitations, recipient and release version	Defined decision or communication	General truth beyond its scope

2.2 Industry information structures

ILPA's Due Diligence Questionnaire provides a standardised framework for manager diligence and was updated in 2021 [1]. Its categories can guide a question taxonomy covering organisation, strategy, team, governance, operations, risk, service providers and responsible investment. ILPA's Reporting Template 2.0, Performance Template guidance, Principles 3.0 and portfolio-company metrics resources provide related structures for fees, expenses, performance, governance and portfolio reporting [2-5]. These resources are industry guidance and templates. They do not determine the completeness, truth or legal effect of a particular manager's documents.

The taxonomy should connect every question to its governing document families. A key-person question may require the limited partnership agreement, private-placement memorandum, side letter and team schedule. A performance question may require audited financial statements, fund cash flows, portfolio-level support, valuation

policy and a stated methodology. A fee question may require governing documents, side letters, capital-account statements and calculations. A single retrieved paragraph rarely supplies the complete answer.

Diligence domain	Typical source families	Evidence-control emphasis
Organisation and team	DDQ, organisation chart, biographies, employment or governance material	Named entity, role, date and current-status reconciliation
Strategy and mandate	PPM, LPA, DDQ, investment policy and track record	Definitions, exclusions and vehicle-specific scope
Performance	Audited statements, cash-flow file, portfolio schedule and methodology	Fund/vehicle identity, gross/net basis, currency, period and formula
Fees and expenses	LPA, side letters, statements, reporting template and invoices	Precedence, investor-specific terms and allocation basis
Governance and conflicts	LPA, committee terms, policies, minutes and disclosures	Authority, approval, conflict and exception history
Operations and controls	Policies, service-provider reports, business-continuity and cyber material	Effective date, scope, exceptions and independent assurance
Portfolio evidence	Portfolio schedule, valuations, KPI reports and board material	Company, date, ownership, source and consistency
Regulatory status	Form ADV or equivalent filing, regulator register and disclosures	Jurisdiction, entity, filing date and current applicability

2.3 Evidence hierarchy and precedence

Authority is claim-specific. An executed agreement can govern a contractual term. An audited financial statement can provide financial evidence within its scope and period. A regulatory filing can evidence what was filed on a date. A DDQ can state management's response. A marketing deck can communicate a claim and may be unsuitable as the sole source for it. The system should store source type and authority without assigning universal rank.

Version precedence requires rules. Executed amendments may supersede earlier terms. Side letters may apply to one investor. Final audited numbers may replace draft accounts. A corrected portfolio schedule may supersede a prior upload. File names and folder dates are weak signals. Precedence should be confirmed from document content, execution status, issuer communication and qualified review.

Evidence tier	Example	Initial treatment	Release gate
D0 unknown	Orphan file, screenshot, missing origin or unreadable scan	Quarantine	Origin, integrity and access resolved
D1 exploratory	Draft deck, analyst note, unverified management response	Hypothesis and request list	Corroborated or labelled representation
D2 controlled representation	Current DDQ, signed response or issued report	Candidate evidence	Scope, date and conflicts reviewed
D3 governing or independently assured evidence	Executed agreement, audited report or official filing	High-authority evidence for its stated purpose	Applicability and qualifications reviewed
D4 accepted decision packet	Reconciled sources, calculations, conflicts and named approval	Defined downstream use	Release, retention and correction controls pass

2.4 Decision boundaries

Document intelligence supports four layers. The source layer preserves files and permissions. The evidence layer extracts and reconciles facts. The professional layer applies legal, accounting, regulatory, tax, valuation, audit or cybersecurity judgement. The investment layer integrates approved evidence with mandate, risk appetite, portfolio construction and committee authority. The system can prepare work for the professional and investment layers. It should not assume their authority.

Question	AI-supported contribution	Required authority
What is the stated management fee?	Retrieve clauses, identify versions, extract rate/basis and expose conflicts	Legal or fund-operations review confirms applicable term

Question	AI-supported contribution	Required authority
What performance is reported?	Link cash flows, statements and reported figures; reproduce approved calculations	Finance, valuation and investment reviewers confirm basis
Is a disclosure complete?	Compare required taxonomy with available evidence and list gaps	Compliance or legal owner determines sufficiency
Is a cyber control effective?	Organise policy, assurance, test and exception evidence	Cybersecurity specialist evaluates design and operation
Does the opportunity fit the mandate?	Map accepted evidence to approved criteria and highlight exceptions	CIO or investment committee decides
May a statement be sent to investors?	Assemble source-linked draft and required legends	Authorised legal/compliance and communication owners approve

2.5 The accountable work unit

The accepted decision-evidence packet prevents productivity measurement from rewarding low-quality output. One packet can answer a single material question or a tightly related set. It should be small enough for a reviewer to challenge and complete enough for that review to be meaningful.

Figure 2. Anatomy of a decision-evidence packet

Source: Matchpoint framework; records and governance context [20-27].

Packet status should move through defined states: received, classified, candidate, conflicted, unanswerable, accepted, released, corrected or retired. Each transition has an owner and timestamp. A draft narrative without the underlying packet cannot enter a controlled investment paper or external answer.

3. DATA-ROOM EVIDENCE MODEL AND SOURCE HIERARCHY

3.1 Ingestion and immutable identity

Ingestion should preserve the original bytes before any conversion. The system records a cryptographic hash, file name, media type, source channel, uploader, receipt time, folder path, access classification and malware-scan result. A derivative PDF, OCR text layer, page image or spreadsheet export receives a separate identity linked to the preserved original. Deduplication should recognise byte-identical files and near-duplicates while retaining provenance.

Records-management principles are relevant. ISO 15489-1 addresses the creation, capture and management of records, including metadata and controls [20]. ISO 18829 describes a framework for trustworthy enterprise content

and electronic document or records management systems [21]. These standards support authenticity, reliability, integrity and usability as design goals. A project can apply their principles without claiming certification or formal conformance.

Ingestion control	Evidence recorded	Failure prevented or exposed
Original preservation	Hash, immutable location and retention class	Silent overwrite or conversion loss
Source identity	Uploader, channel, received time and declared issuer	Orphan evidence
File inspection	Media signature, encryption, malware and corruption result	Unsafe or unreadable processing
Duplicate analysis	Exact hash, perceptual/text similarity and document family	Double counting and conflicting duplicates
Access label	Matter, investor, vehicle, confidentiality and permitted purpose	Cross-matter disclosure
Conversion lineage	Tool, version, parameters and output hash	Unreproducible OCR or rendering
Catalogue status	Active, superseded, quarantined or retired	Retrieval from invalid source

3.2 Document families and entity resolution

A data room contains families rather than isolated files. A PPM may have supplements. An LPA may have amendments. Side letters apply to named investors. Financial statements recur by period. A portfolio schedule may be revised without a clear name. The catalogue should group likely families, then require reviewer confirmation where precedence matters.

Entity resolution is equally important. A manager brand may contain advisers, general partners, carry vehicles, funds, parallel vehicles, feeders and portfolio companies. The same abbreviated name can refer to different entities. Every candidate fact should carry the legal or operational entity to which it applies. Where that identity is ambiguous, the packet remains conflicted or unanswerable.

3.3 Page, paragraph, table and cell coordinates

Evidence must be locatable. For PDF material, page number alone can be insufficient when printed and electronic numbers differ. The evidence span can store the PDF page index, printed label, bounding box, extracted text and a page-image hash. For spreadsheets, it can store workbook version, sheet name, cell range, formula, displayed value and named-range identity. For structured filings, it can store taxonomy concept, context, unit and fact identifier.

The SEC describes Inline XBRL as a format that combines human-readable disclosure and machine-readable data [11]. The SEC publishes investment-company taxonomies for structured filings [12]. The IFRS Foundation publishes the IFRS Accounting Taxonomy, and the 2025 taxonomy remains current for 2026 [38]. XBRL International's Inline XBRL and core specifications define the related technical structures [39,40]. Structured facts can reduce extraction ambiguity; their contexts, units, extensions and filing scope still require review.

3.4 Access follows the evidence

Permission filtering should occur before retrieval. A user who cannot open a source should not receive its extracted text, embedding, summary, answer, citation, cached prompt or log content. Matter, vehicle, investor and confidentiality labels should propagate to derivatives. Retrieval results should be intersected with the caller's current entitlements at query time.

NIST SP 800-207 frames zero trust around protecting resources and granting access through explicit policy rather than relying on network location [23]. ISO/IEC 27001:2022 specifies requirements for an information-security management system [22]. These sources support identity, least privilege, asset protection and monitored access. They do not prescribe one product architecture.

Access layer	Required control	Test evidence
Source vault	Matter-scoped roles, encryption and immutable audit	User cannot list or fetch an unauthorised original
OCR and parsed store	Inherited source label and segregated keys	Parsed text remains unavailable without source entitlement
Vector or search index	Per-chunk access attributes and query-time filtering	Cross-matter retrieval test returns zero results

Access layer	Required control	Test evidence
Model context	Approved endpoint, no unauthorised retention/training and redaction rules	Provider settings and controlled prompt-log sample
Output store	Recipient, purpose, expiry, watermark or export rule	Download and onward-share tests
Logs and telemetry	Minimized content, protected identifiers and retention schedule	Operations staff cannot read deal content without authority

3.5 Source conflict is a first-class object

Conflicts should not be resolved by whichever document ranks highest in similarity search. The system records the conflicting propositions, each evidence span, document dates, entity scope, governing status and assigned owner. Deterministic rules can resolve simple cases, such as a confirmed final version replacing a draft. Material contractual, accounting, valuation or regulatory conflicts require qualified review.

Examples include gross performance in a deck versus net performance in audited statements; a DDQ team count versus a current organisation chart; a PPM fee description versus an investor-specific side letter; and a portfolio schedule valuation versus a later audited figure. The released packet should show which source was accepted, why, by whom and for what use.

4. A2 AND B2 USE CASES

4.1 A2 manager screening and diligence planning

A family-office CIO can use document intelligence to organise an initial manager screen. The system can map available documents to an ILPA-informed question taxonomy, identify missing periods or governing documents, extract candidate organisational and fund facts and prepare an evidence-linked request list. The output is a diligence plan, not an investment score.

The screen should distinguish management representations from independently assured or governing evidence. It should also distinguish absence from failure. A document may be unavailable because the diligence stage is early, because access is restricted or because the manager has not supplied it. The packet records the reason where known and labels it unknown otherwise.

4.2 A2 track-record and portfolio review

Track-record review is highly context-dependent. The system can ingest controlled cash-flow data, reported performance, portfolio schedules and valuation policies; it can map each input to fund, vehicle, currency, date and gross-or-net basis. Deterministic code can reproduce approved calculations. Investment staff then assess persistence, attribution, comparability, valuation practice, concentration and relevance to the proposed strategy.

Document intelligence should not infer missing cash flows or normalise different strategies without an approved methodology. It should expose excluded deals, partial realisations, subscription facilities, foreign-exchange treatment, recycled capital, fees and carry where those factors affect the stated metric. The calculation record contains inputs, formula, code version and reviewer.

A2 use case	Candidate automation	Acceptance requirement	Named owner
Data-room inventory	Classify files and map them to the question taxonomy	Reviewer confirms families, versions and gaps	Diligence lead
Organisation review	Extract people, roles, dates and changes	Current source and entity reconciliation	Investment team
Track-record preparation	Link cash flows, portfolio facts and reported metrics	Approved data and reproducible calculation	Finance/investment reviewer
Terms comparison	Retrieve and compare relevant clauses	Governing and investor-specific applicability confirmed	Legal reviewer
Operational diligence	Assemble controls, service-provider and exception evidence	Specialist review of design and operation	Operations/cyber owner
IC paper support	Draft source-linked factual sections and open questions	Every material claim maps to accepted packet	Deal lead and CIO

4.3 A2 investment-committee preparation

An investment paper benefits from claim-level lineage. Facts about fund size, target return, team, strategy, fees, governance and performance should link to accepted packets. Material conflicts and unanswered questions should appear alongside the draft. The system can produce an evidence map and change log when new documents arrive.

The committee should see decision-relevant uncertainty. A missing side letter, pending audit, unresolved key-person definition or inconsistent portfolio valuation can be more important than dozens of extracted fields. The workflow should prioritise material exceptions over document volume.

4.4 B2 fundraising data-room readiness

A GCC fund manager can use the framework before granting investor access. The system can catalogue the room, identify duplicate or superseded files, check taxonomy coverage, reconcile recurring facts and flag stale dates. It can prepare a controlled source register and owner-assigned remediation list.

Readiness is not completeness by file count. The manager should confirm that statements across the PPM, DDQ, pitch deck, track-record workbook, website and regulatory filings are consistent or properly qualified. Each release has an approved source and date. Investor-specific material remains segregated.

4.5 B2 investor-question response

Investor questionnaires repeat themes while differing in wording and scope. Retrieval can propose prior approved evidence and draft a response. The workflow should verify the requesting investor, vehicle, current period, permitted disclosure and applicable side-letter or confidentiality context. Previous answers are reference material rather than automatic truth.

The response packet should include the requested question, proposed answer, evidence spans, differences from the last approved answer, owner and expiry or review date. A human approves every external response. Questions requesting legal interpretations, regulatory representations, performance claims or forecasts receive specialist review.

4.6 B2 marketing and disclosure control

The SEC's investment-adviser marketing rule framework addresses advertisements, testimonials, endorsements, performance and related recordkeeping [6]. SEC staff maintains marketing-compliance FAQs and states that staff statements lack legal force [7]. Current applicability depends on adviser status, jurisdiction, communication and facts. A document-intelligence system can connect a proposed claim to substantiation, approval and retained evidence. Legal and compliance owners determine whether it may be used.

The system should distinguish actual, hypothetical, extracted and target performance; gross and net; realised and unrealised; fund and deal; and current and superseded figures. It should preserve qualifications and prevent a draft from entering an external channel without approval. An approved content library can reduce repeated work when every item has an evidence set, audience, use limitation and expiry date.

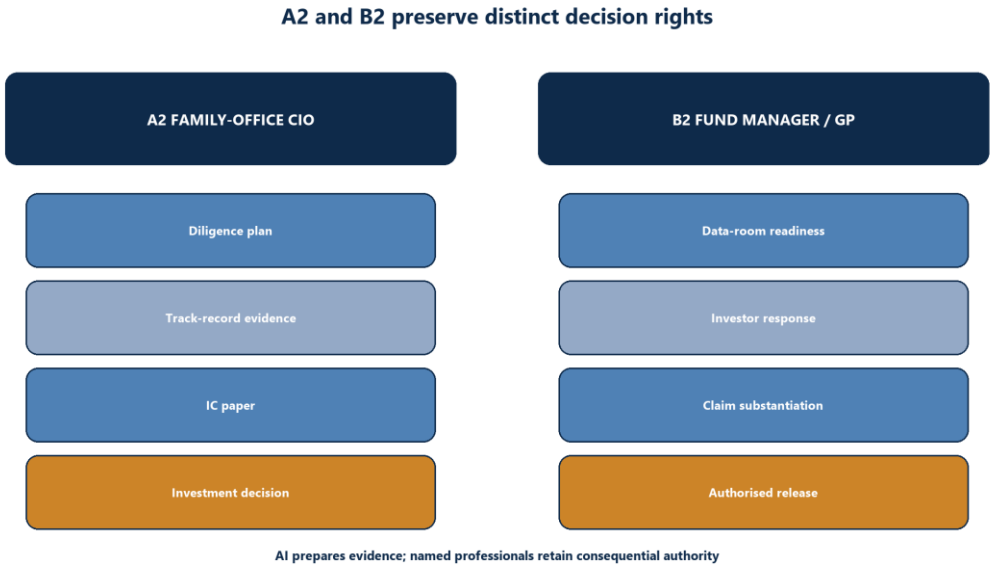

Figure 3. A2 and B2 decision-rights map

Source: Matchpoint operating framework.

4.7 Shared stop conditions

Stop condition	Required response	Owner
Governing version unresolved	Hold answer and route document family for precedence review	Legal/fund operations
Entity, vehicle, currency, unit or period ambiguous	Mark candidate fact unanswerable	Evidence reviewer
Source access does not permit recipient or purpose	Block retrieval and release	Information owner
OCR or table structure fails a material field	Reprocess or transcribe under dual review	Data owner
Calculation cannot be reproduced	Hold metric and repair lineage	Finance/model owner
External statement lacks substantiation or approval	Block publication or response	Compliance/legal
Model instruction is found inside a source document	Treat as untrusted content and isolate	Security/model owner
Reviewer capacity or exception queue exceeds limit	Pause scale and restore manual priority workflow	Process owner

5. DOCUMENT-INTELLIGENCE METHODS AND FAILURE MODES

5.1 OCR and document rendering

OCR converts page images into text and coordinates. Digitally generated PDFs may already contain text, yet extraction order can be wrong when the page has columns, footnotes, floating text boxes or complex tables. Scans add rotation, noise, compression, handwriting, stamps and faint print. The workflow should preserve the page image and compare extracted text with the visual source.

OCR quality should be evaluated on the fields that matter. A low character-error rate can coexist with a material error in a fee rate, date, currency or negative sign. Validation should include critical-field exact match, numerical consistency and human review of low-confidence or high-consequence spans.

5.2 Layout-aware representation

LayoutLM introduced joint modelling of text and document layout for document image understanding [28]. LayoutLMv2 added text, layout and image interaction in a multimodal pretraining framework [29]. DocLLM uses bounding-box information to support reasoning over visually rich documents [32]. These studies establish that

spatial layout can improve selected document tasks. They do not establish performance on a private-market data room or remove the need for task-specific validation.

Layout matters because a number can inherit meaning from a row label, column header, footnote and table title. A model that extracts visible tokens without these relationships can assign the wrong period, entity or unit. The evidence span should retain its structural neighbourhood.

5.3 Table extraction

Private-market evidence often resides in tables: cash flows, portfolio schedules, fee tables, organisation charts, KPI grids and financial statements. Table extraction should preserve rows, columns, merged cells, headers, notes and repeated page headers. A visually split table may need to be reconstructed across pages. A spreadsheet requires formula and displayed-value handling.

Table failure	Example consequence	Control
Header drift	2024 value assigned to 2025	Header-path preservation and row/column validation
Merged-cell loss	Vehicle or currency scope disappears	Structural reconstruction and visual overlay
Footnote detachment	Adjusted metric presented without qualification	Footnote-to-cell linkage
Sign or decimal error	Loss becomes gain or scale changes	Format-aware parsing and arithmetic checks
Page split	Portfolio rows omitted or duplicated	Table-family stitching with page evidence
Formula/value confusion	Stale displayed value treated as recalculated result	Store formula, cached value and calculation status
Hidden rows/columns	Excluded data omitted from review	Workbook-structure inventory and disclosure

5.4 Document visual question answering

DocVQA assembled more than 50,000 questions over more than 12,000 document images and reported a substantial gap between human and model results, especially on questions requiring structural understanding [30]. DocCVQA extended the task to questions over collections of documents [31]. These benchmarks are relevant because diligence questions often require cross-document aggregation. Their datasets and metrics remain different from legal, financial and investor-document decisions.

Question answering should return evidence and an answerability status. SQuAD 2.0 deliberately combines answerable and unanswerable questions, making abstention part of the task [34]. A diligence system needs stronger conditions: evidence can be absent, inaccessible, contradictory, stale, entity-mismatched or outside professional authority. The output should say which condition applies.

5.5 OCR-free and multimodal generation

Donut proposed an OCR-free document understanding transformer to avoid error propagation from a separate OCR engine [33]. Multimodal models can reason across page images and text. They still require source preservation, permissions, deterministic field checks, version control and validation. An end-to-end answer that cannot expose stable evidence coordinates is unsuitable for material release.

5.6 Retrieval and generation

Retrieval-augmented generation can select relevant chunks and place them in a model context. Retrieval quality depends on parsing, chunking, metadata, embeddings, query formulation, access filtering and ranking. Generation quality depends on context, instructions, model version and decoding. Each layer can fail independently.

The workflow should retain the complete retrieval set and the evidence actually cited. It should report when a relevant document was not indexed, a chunk lost a table header, an access rule removed a source or the answer exceeded the supported evidence. A citation to a nearby page is not proof that the sentence follows from it.

5.7 Prompt injection and untrusted document content

A source document can contain text that attempts to instruct the model to ignore policy, reveal data or call a tool. NIST's adversarial machine-learning taxonomy covers evasion, poisoning, privacy and misuse risks across AI systems and includes generative-AI attack concepts [26]. The NIST Generative AI Profile also identifies risks and

risk-management actions for generative systems [25]. Data-room content should be treated as untrusted evidence, never as operating instructions.

Defences include separating system policy from content, disabling unnecessary tool authority, sanitising active content, enforcing retrieval permissions outside the model, constraining output schemas, monitoring anomalous instructions and using human approval for consequential actions. A model should not send emails, alter a data room, approve a response or publish a claim merely because a document requests it.

5.8 Failure-mode register

Failure mode	Observable symptom	Control and test
Wrong document family	Answer cites a draft or unrelated vehicle	Family/version test set and precedence review
OCR corruption	Material number or negation differs from page	Critical-field exact match and visual overlay
Chunk-context loss	Rate loses basis, period or exception	Structural chunking and context completeness test
Retrieval miss	Relevant source exists but is absent from evidence set	Recall-at-k on answerable reference questions
Unsupported generation	Claim lacks entailment from cited spans	Claim-level evidence review and abstention
Cross-document contradiction	One fluent answer hides divergent sources	Conflict detector and multi-source packet
Access leakage	Restricted fact appears in another user's result	Adversarial permission test across every store
Prompt injection	Source content changes system behaviour	Untrusted-content harness and restricted tools
Numerical error	Units, signs, periods or formulas are wrong	Typed schema, deterministic calculation and reconciliation
Model/version drift	Acceptance or error mix changes after release	Locked regression set and release gate

6. RETRIEVAL, EXTRACTION, CALCULATION AND CONFLICT RESOLUTION

6.1 Question contracts

Each material question should have a contract. The contract defines the entity, vehicle, period, currency, unit, acceptable source families, minimum authority, output schema, reviewer, answerability conditions and prohibited inference. A question such as “What is the management fee?” is incomplete until it identifies the fund, investor class, investment period, basis and applicable agreements.

The contract also defines evidence sufficiency. A key-person answer may require both the governing definition and current team facts. A performance answer may require cash flows and methodology. A regulatory-status answer may require an official register check in addition to supplied documents. The system returns incomplete when the required combination is absent.

6.2 Retrieval tests

Retrieval should be measured on a reference set built from representative documents and questions. Metrics can include whether all necessary evidence spans appear within the reviewed set, whether prohibited sources are excluded and whether conflicts are surfaced. Precision matters for reviewer load; recall matters for missed evidence. The relevant operating measure is accepted packets per reviewer hour at an approved critical-error level.

Retrieval test	Unit	Pass logic
Answerable question recall	Required evidence spans found within review set	All material spans present for the question contract
Conflict recall	Known conflicting proposition sets surfaced	Every seeded material conflict routed
Version precision	Retrieved sources use approved precedence	No superseded source presented as current
Access precision	Retrieved chunks permitted for caller and purpose	Zero unauthorised chunks in adversarial test
Citation fidelity	Citation coordinates resolve to supporting content	Reviewer can reproduce every material claim
Abstention	Unanswerable cases withheld with correct reason	No forced answer in seeded absence/ambiguity cases

6.3 Candidate-fact schema

A candidate fact should be typed. A numerical field includes value, sign, unit, scale, currency, period start/end, as-of date, entity, vehicle, gross/net or actual/forecast basis, source coordinates and extraction confidence. A clause field includes defined term, operative text, exceptions, governing document, version and applicable party. A person field includes name, role, entity, effective date and source.

Confidence is diagnostic rather than authority. A model may be highly confident and wrong. A low-confidence span can be correct. Acceptance combines source authority, extraction validation, reconciliation and reviewer judgement.

6.4 Deterministic calculation service

Material calculations should occur outside free-form generation. The calculation service accepts approved typed facts, checks units and periods, applies versioned formulas and returns the result with an execution record. The language model can explain an approved result and cite its inputs. It should not silently calculate from prose or fill missing values.

Performance calculations require a documented methodology. The ILPA performance guidance provides industry context for standardised reporting [3]. Applicable accounting, valuation and investor-reporting requirements remain case-specific. A calculation should state cash-flow conventions, dates, currency treatment and adjustments; the reviewer approves its use.

6.5 Conflict handling

Conflicts move through triage, classification, assignment and resolution. The system determines whether the difference is exact, rounding, period, entity, definition, version or substantive. It can resolve deterministic cases under approved rules. It cannot select a legal interpretation or preferred valuation merely because one source is newer.

Conflict class	Example	System action	Human authority
Version	Draft DDQ differs from approved DDQ	Apply confirmed supersession; retain both	Document owner confirms
Scope	Fund figure differs from strategy figure	Preserve both with entity/vehicle tags	Investment/finance reviewer
Basis	Gross differs from net performance	Label bases and block comparison until aligned	Finance/investment reviewer
Contract	LPA differs from investor side letter	Route applicable party and clause set	Legal reviewer
Timing	Portfolio value differs by reporting date	Preserve as-of dates and change history	Valuation/finance reviewer
Representation	Deck claim differs from audited support	Flag substantiation gap and block external reuse	Compliance/legal

6.6 Abstention and escalation

The system should abstain with a reason code: missing evidence, inaccessible evidence, unreadable evidence, ambiguous entity, ambiguous period, conflicting sources, unsupported calculation, outside authority or policy restriction. It should propose the next evidence request or named reviewer. An empty answer should not be replaced with a plausible industry norm.

6.7 Change and correction

When a new document arrives, the system identifies affected packets and outputs. A corrected audited figure may invalidate an IC draft, an investor answer and a marketing claim. Dependency links enable targeted review. Released content receives a new version; prior versions remain retained according to policy. The correction record identifies what changed, why, who approved it and which recipients require notification.

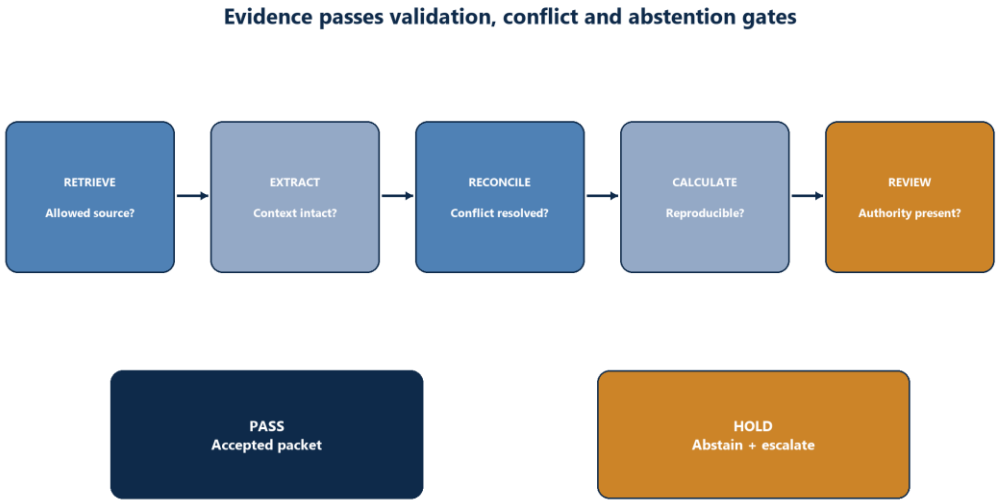

Figure 4. Validation, conflict and abstention gates

Source: Matchpoint framework; unanswerable-question and governance context [24-26,34].

7. CONTROLLED ARCHITECTURE AND TOOL STACK

7.1 Architectural principles

The architecture should enforce controls outside the generative model. Identity, access, source immutability, version state, typed schemas, deterministic calculations, release authority and audit logs are system responsibilities. The model can support classification, extraction, retrieval and drafting within those controls.

NIST AI RMF 1.0 organises AI risk work through Govern, Map, Measure and Manage functions [24]. ISO/IEC 42001:2023 specifies requirements for an AI management system [27]. These frameworks support documented context, accountability, risk assessment, measurement and continual improvement. A project should state which controls it applies rather than claim conformance without evidence.

7.2 Reference architecture

The source vault stores originals and access metadata. A conversion service produces controlled page images, text, layout and spreadsheet structures. The catalogue holds document families, versions, entities and provenance. The evidence store contains spans and candidate facts. A permission-aware search service retrieves evidence. A calculation service handles approved formulas. A review workbench shows source and derivative together. A release registry records accepted packets and downstream uses.

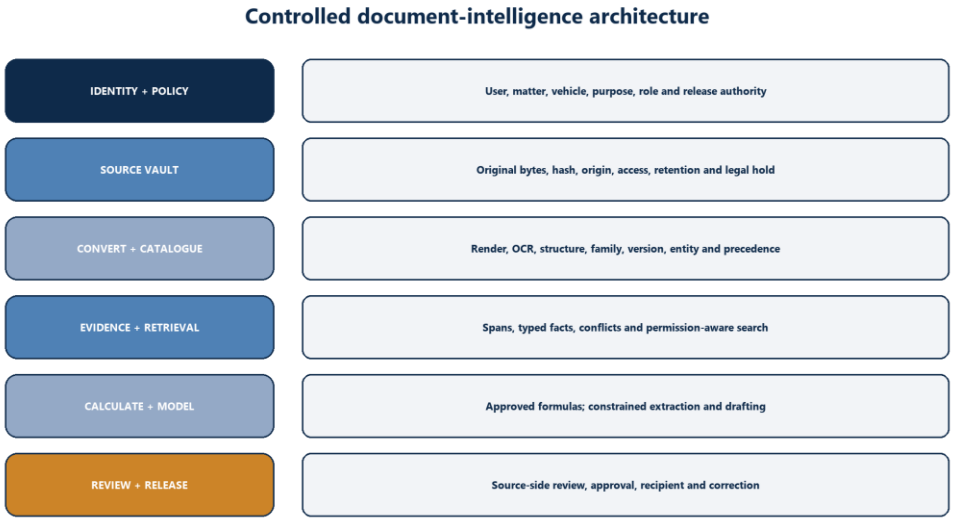

Figure 5. Controlled document-intelligence architecture
Source: Matchpoint architecture; security and governance context [20-27].

Layer	Core responsibility	Release evidence
Identity and policy	User, matter, vehicle, purpose, role and approval authority	Access decision and policy version
Source vault	Original bytes, hash, origin, retention and legal hold	Immutable source record
Conversion	Render, OCR, parse, structure and malware isolation	Tool/version lineage and QA result
Catalogue	Family, version, entity, source type and authority	Reviewer-confirmed precedence
Evidence store	Spans, candidate facts, conflicts and dependencies	Stable source coordinates
Retrieval	Permission-filtered search and question contract	Retrieval set and ranking version
Calculation	Typed inputs and versioned deterministic functions	Reproducible execution record
Model service	Approved classification, extraction and drafting tasks	Model, prompt, context and output version
Review workbench	Side-by-side evidence, exception and approval	Named disposition and comments
Release registry	Recipient, purpose, version, expiry and correction	Signed release record

7.3 Model and vendor boundary

The vendor assessment should address data use, retention, regional processing, encryption, tenant separation, subprocessors, incident response, model changes, evaluation, deletion, export and audit evidence. A “no training” statement should be confirmed contractually and technically for the selected service configuration. Provider marketing material is not sufficient evidence of the deployed control state.

Models and prompts are versioned components. A model update can change extraction and abstention behaviour. Production changes should run through a locked regression set, security tests and reviewer acceptance. Rollback should be practical.

7.4 Observability without content overcollection

Operations need latency, error, retrieval, acceptance, override and security telemetry. Logs should minimise sensitive content. Identifiers can link an event to protected evidence without copying full document text into general monitoring systems. Access to logs follows defined roles and retention.

7.5 Separation of environments

Development, validation and production should be separated. Synthetic or appropriately controlled test documents can support development. Representative validation requires authorised documents and a defined protocol. Production data should not move into personal tools or uncontrolled sandboxes. Export and local-download rules should match matter policy.

7.6 Resilience and manual fallback

The process should function when a model, index or vendor is unavailable. The source catalogue, request list, high-priority evidence and manual review path remain accessible. Recovery tests should cover source restoration, index rebuild, audit-log integrity and pending-review reconciliation. Business continuity should prioritise material deadlines and investor commitments.

8. VALIDATION, QA AND HUMAN AUTHORITY

8.1 Representative validation corpus

The validation corpus should reflect the documents and questions intended for use. It should include scans, native PDFs, spreadsheets, long agreements, tables across pages, amendments, side letters, repeated questions, multiple vehicles, currencies, date formats and deliberately unanswerable cases. Development and test sets should be separated by document family or fund where leakage would inflate results.

The corpus should include critical fields and adverse cases. Rare but consequential errors, such as a wrong fee, sign, currency, key-person condition or performance basis, deserve explicit thresholds. Aggregate token or field accuracy can hide them.

8.2 Quality gates

Gate	Evidence	Illustrative pilot disposition
Source integrity	Original retrievable; hash and origin match	Stop on any unresolved material source
Access isolation	Adversarial cross-matter tests	Stop on any unauthorised disclosure
Critical-field extraction	Exact match for approved critical schema	Threshold requires CK approval and pilot evidence
Retrieval completeness	Required spans within reviewer set	Threshold requires CK approval and representative test
Conflict detection	Seeded material conflicts surfaced	Stop on any missed critical conflict
Abstention	Unanswerable and outside-authority cases withheld	Stop on forced material answers
Calculation	Reproduction from approved inputs	Zero unexplained difference
Citation fidelity	Claim maps to stable supporting coordinates	Stop on unsupported material claim
Reviewer workflow	Queue, override and release recorded	Named owner and service-level rule approved
Regression	Locked corpus before every model/prompt change	No unapproved critical regression

All threshold language in this paper is a proposed framework for management approval. No Matchpoint-approved T15 threshold was supplied.

8.3 Human review levels

Review depth should follow consequence and reversibility. Low-consequence classification can use sampled review after validation. Material financial, contractual, regulatory and external claims receive full review. The reviewer should see the original page or workbook context, not only extracted text. Override reasons become validation evidence.

Review level	Example task	Review rule	Release authority
R0 exploratory	Internal search or taxonomy suggestion	No downstream release	Analyst
R1 administrative	File classification after validated pilot	Sampling plus exception review	Data-room owner
R2 factual	Team fact or issued-date extraction	Source-side review	Diligence reviewer

Review level	Example task	Review rule	Release authority
R3 material numerical	Fee, performance, valuation or exposure fact	Full evidence and calculation review	Finance/investment owner
R4 contractual/regulatory	Governing term or regulated statement	Qualified full review	Legal/compliance owner
R5 investment/external	IC recommendation or investor communication	Approved packets plus mandate review	CIO/IC or authorised signatory

8.4 Reviewer calibration

Human review is not automatically consistent. A sample should be independently double-reviewed to measure agreement and refine definitions. Disagreements reveal ambiguous questions, weak source hierarchy or inadequate training. Reviewers need escalation paths and protected time for material exceptions.

8.5 Monitoring in production

Monitoring should include document mix, answerability, acceptance, correction, override, queue age, latency, access denials and critical incidents. Changes by fund, document type, model version and reviewer can reveal drift. Acceptance rate alone is unsafe because reviewers can accept weak outputs under time pressure.

8.6 Red-team and security testing

Tests should include malicious document instructions, hidden text, active links, oversized files, unusual encodings, cross-matter query attempts, data exfiltration prompts, poisoned indexes and unauthorised export. NIST's generative-AI and adversarial-ML sources provide risk context [25,26]. Security owners should define the actual test plan and remediation evidence.

9. PRODUCTIVITY, ECONOMICS AND COMMERCIAL ATTRIBUTION

9.1 Evidence boundary for productivity claims

Noy and Zhang studied 453 professionals completing selected writing tasks and reported average reductions in time and improvements in output quality for participants using a generative AI tool [35]. Brynjolfsson, Li and Raymond studied 5,179 customer-support agents and reported an average productivity increase measured as issues resolved per hour [36]. Dell'Acqua and co-authors found task-dependent effects around an uneven capability frontier in a field experiment with consultants [37]. These studies concern their own tasks, populations, tools and periods. They do not establish a productivity or ROI result for private-market data-room diligence.

T15 therefore uses a pilot model rather than a forecast claim. The baseline and assisted workflow should be observed on comparable packets. The primary numerator is accepted packets. The denominator is total reviewer and operator time. Critical-error, correction and exception measures remain co-primary gates.

9.2 Illustrative workflow model

The following scenario is an unverified illustrative management assumption. It models 120 decision-evidence packets in one month. Baseline preparation is assumed at 1.00 hour per packet and review at 0.50 hour. The assisted scenario assumes 0.35 hour of preparation, 0.45 hour of review, 18 hours of monthly implementation/operations and 12 hours of rework. These figures were not observed from Matchpoint work.

Scenario component	Baseline hours	Assisted hours	Status
Preparation for 120 packets	120	42	Unverified illustrative assumption
Review for 120 packets	60	54	Unverified illustrative assumption
Implementation and operations	0	18	Unverified illustrative assumption
Rework and exceptions	12	12	Unverified illustrative assumption
Total	192	126	Illustrative arithmetic only
Accepted packets	120 assumed	120 assumed	Acceptance and critical-error gates must pass
Hours per accepted packet	1.60	1.05	Illustrative arithmetic only

The arithmetic implies 66 hours of gross illustrative time difference and 0.55 hours per accepted packet. It is not a forecast, realised saving or client result. A pilot must replace every input with observed data and include error, queue and reviewer effects.

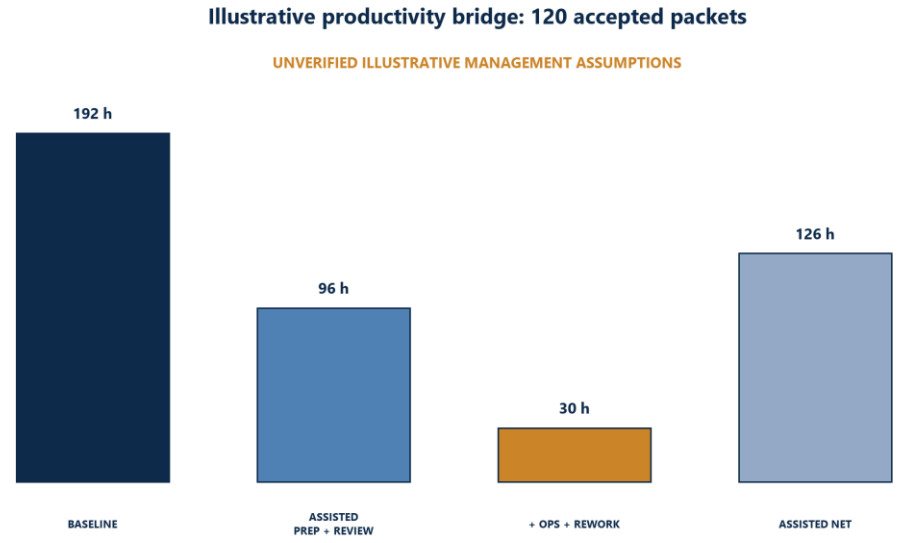

Figure 6. Illustrative productivity and cost bridge

Source: Unverified illustrative management assumptions; Matchpoint calculation.

9.3 Cost and break-even model

Let (Q) be accepted packets, (H_b) baseline hours per accepted packet, (H_a) assisted hours per accepted packet, (C_h) fully loaded hourly cost, (C_f) fixed monthly implementation cost and (C_v) variable system cost per packet. An illustrative monthly operating difference is:

\[\Delta C = Q(H_b-H_a)C_h - C_f - QC_v. \]

Every input requires management approval and observed evidence. Quality gates are constraints rather than monetised benefits. A project should not assign value to an avoided regulatory, legal or investment loss without approved causal evidence.

Economic field	T15 current value	Required evidence for update
Attributed revenue	USD 0	Approved transaction/client record and causal attribution
Attributed cost reduction	USD 0	Observed comparable baseline, assisted cost and approval
Attributed loss reduction	USD 0	Approved incident/counterfactual methodology and evidence
Approved implementation budget	USD 0	Management-approved budget
Observed break-even volume	Not available	Approved costs and representative operating data
Validated productivity uplift	Not available	Completed pilot passing quality and security gates

9.4 Capacity is not automatically revenue

Time released can reduce backlog, improve diligence depth, increase response capacity or remain unused. Revenue requires a separate commercial mechanism and attribution record. For A2, potential value may arise through more timely and consistent decision preparation. For B2, it may arise through faster controlled investor responses and stronger data-room readiness. Neither outcome has been observed or attributed for T15.

9.5 Pilot measurement plan

The pilot should sample representative packet types, stratified by consequence and document complexity. Baseline cases use the current workflow. Assisted cases use the controlled system with equivalent review. The study records total elapsed time, active preparation time, review time, acceptance, correction, critical errors, retrieval completeness and reviewer experience. Repeated or paired cases require leakage controls.

Measure	Definition	Evidence source
Accepted packets per reviewer hour	Packets meeting all release gates divided by review time	Workflow and release registry
Critical unsupported-claim rate	Material claims without adequate evidence per reviewed packet	Independent QA sample
Critical extraction-error rate	Wrong material field per validated field	Locked validation corpus
Conflict recall	Known material conflicts surfaced	Seeded and observed conflict set
Correction rate	Released packets corrected within defined period	Release registry
Queue age	Time from candidate to reviewer disposition	Workflow timestamps
Access incident rate	Unauthorised content exposure events	Security logs and incident process
Total cost per accepted packet	Approved people, vendor and operating costs divided by accepted packets	Finance-approved cost model

10. LEGAL, REGULATORY, PRIVACY AND RECORDS BOUNDARIES

10.1 Investment-adviser records and communications

SEC Rule 204-2 specifies books-and-records requirements for registered investment advisers [8]. The SEC's electronic-recordkeeping rule addresses electronic copies, safeguards, access limitation, completeness, truth and legibility [9]. Form ADV and IARD materials provide filing context [10]. Applicability depends on entity, registration, record and jurisdiction. Counsel and compliance owners should map the actual workflow.

The system should retain source and released-answer histories according to approved schedules, legal holds and correction procedures. Generative prompt logs can themselves contain regulated or confidential records. Record status should be determined by policy, not by whether content sits inside an AI service.

10.2 UAE and financial-free-zone context

The UAE Personal Data Protection Law, Federal Decree-Law No. 45 of 2021, establishes a federal personal-data regime with scope and exceptions defined in the law [16]. ADGM's Data Protection Regulations 2021 apply within their stated scope [17]. DIFC and other jurisdictions have their own regimes. A manager or investor should identify controller/processor roles, lawful basis, notice, data-subject rights, transfer, security, retention and breach obligations for the actual entities and processing.

DFSA materials describe collective-investment-fund categories and fund-document context in the DIFC [13]. ADGM FSRA publishes its regulations and rules and a getting-started guide that identifies FUND Rules as a principal reference for funds [14,15]. These sources establish regulatory context. They do not determine an individual firm's permission, disclosure sufficiency or compliance.

10.3 GDPR and EU AI Act

The GDPR applies according to its territorial and material scope and sets requirements for processing personal data [18]. A data room may contain biographies, contact information, ownership details, compensation data, background checks and other personal data. The system should minimise collection, restrict access and document retention and transfer decisions.

Regulation (EU) 2024/1689, the EU AI Act, entered into force in 2024 and uses staged application dates [19]. The paper's evidence cut-off is 1 August 2026; most provisions are scheduled to apply from 2 August 2026, subject to the regulation's transitional structure. Applicability to a particular provider, deployer, system or use requires qualified analysis. The architecture should maintain system inventory, purpose, provider information, risk assessment, instructions, monitoring and human-oversight evidence that can support such analysis.

10.4 Confidentiality, privilege and contractual restriction

Data-room access agreements, non-disclosure agreements, engagement terms and fund documents can restrict use, recipients, copying, model training, retention and onward disclosure. Privileged material may require segregation. The system should capture permitted purpose and recipient at source ingestion and release. Legal review determines privilege and contractual effect.

10.5 Marketing substantiation

Any claim about performance, team, assets, strategy, ESG, operations or AI-enabled productivity needs appropriate support and approval. The SEC marketing materials are relevant for US-registered investment advisers [6,7]. Other jurisdictions and communication types have separate rules. The release registry should connect each external statement to its evidence, audience, approval, qualifications and expiry.

10.6 Professional boundaries

Document intelligence does not provide legal interpretation, audit assurance, accounting treatment, tax advice, valuation opinion, regulatory approval, cybersecurity assurance or investment recommendation. It can organise materials for qualified professionals. Outputs should state their evidence basis and limitations.

Boundary	System may support	System must route
Legal	Clause retrieval, version map and issue list	Applicability, interpretation, privilege and advice
Accounting/audit	Evidence organisation and reproducible approved calculation	Accounting treatment and assurance conclusion
Valuation	Source reconciliation and exception analysis	Method, assumptions and valuation opinion
Compliance	Substantiation packet and approval workflow	Regulatory interpretation and permission
Cybersecurity	Policy/evidence catalogue and control-question mapping	Effectiveness assessment and assurance
Investment	Source-linked IC paper and open-question register	Recommendation and committee decision

11. A2 AND B2 OPERATING PLAYBOOKS

11.1 A2 family-office playbook

The A2 charter begins with one strategy or manager and a defined committee decision. The CIO names the diligence lead, professional reviewers and release authority. The team maps the existing question set to required source families and selects a representative corpus. It excludes personal, privileged or restricted material until permissions are confirmed.

The initial pilot should focus on evidence organisation, question coverage and factual packet preparation. Contract interpretation, performance conclusions and recommendations remain fully human-controlled. The committee paper displays accepted evidence, conflicts and unanswered questions. After each meeting, the team records which packets were useful, corrected or superseded.

A2 stage	Activity	Gate	Deliverable
Charter	Define decision, corpus, roles and excluded data	CIO and information-owner approval	Signed pilot charter
Baseline	Measure current packet time, quality and backlog	Representative sample confirmed	Baseline register
Catalogue	Preserve, classify and version documents	Source and access QA pass	Controlled data-room map
Evidence pilot	Retrieve and extract selected question types	Critical quality and abstention gates pass	Shadow packets
IC shadow	Compare assisted and current paper preparation	No unauthorised or unsupported material	Evaluation report
Restricted use	Release approved factual sections with full review	Named sign-off and rollback	Controlled IC workflow
Scale/retire	Expand by tested task or withdraw	Risk-owner decision	Approved operating state

11.2 B2 fund-manager playbook

The B2 charter begins with one fundraise, controlled source register and investor-response process. The manager names document owners across legal, compliance, finance, investment, operations and IR. The catalogue identifies governing documents, approved figures, current narratives and investor-specific restrictions.

The pilot can automate taxonomy mapping, document-family grouping, stale-date checks and draft evidence packets for selected factual questions. External answers remain subject to the existing approval matrix. The release registry records the investor, question, answer, sources, approval and date. New facts trigger review of affected prior answers.

B2 stage	Activity	Gate	Deliverable
Charter	Define fund, investors, question families and authority	GP/compliance approval	Signed use-case charter
Source register	Reconcile PPM, LPA, DDQ, deck, track record and filings	Version owners confirm	Approved source register
Readiness review	Map coverage, conflicts, stale facts and restrictions	Remediation owners assigned	Data-room readiness report
Response shadow	Draft evidence-linked responses without sending	Material claims fully supported	Shadow response packets
Approval integration	Connect legal/compliance/finance review	Existing sign-off matrix preserved	Controlled response queue
Restricted release	Use for defined factual questions and recipients	Recipient/purpose checks pass	Approved investor answers
Scale/retire	Add question classes after validation or withdraw	GP/risk-owner decision	Approved operating state

11.3 Shared decision-rights matrix

Decision	Accountable owner	Required contributors	Evidence retained
Approve use case	CIO/GP or delegated risk owner	Legal, compliance, data and process owners	Charter and risk assessment
Admit source	Information owner	Legal/security where required	Origin, access and retention decision
Confirm version precedence	Document owner	Legal, finance or fund operations	Version disposition
Accept material fact	Domain owner	Evidence reviewer and specialist	Packet and review record
Approve model/prompt release	Model/risk owner	Security, validation and process owners	Regression and security results
Release IC or investor output	CIO/IC or authorised signatory	Legal/compliance/finance as applicable	Released version and approvals
Correct or retire output	Original release owner	Affected domain and recipients	Correction and notification record

12. GATED ADOPTION ROADMAP

12.1 Stage 0: charter and evidence boundary

Define the decision, user, corpus, prohibited data, source authority, review roles, success measures and stop conditions. Record that T15's current commercial attribution is USD 0. Select vendors only after information-security, privacy and contractual review.

12.2 Stage 1: baseline and corpus

Observe the current workflow on a representative sample. Build the locked validation corpus and answer key under appropriate permissions. Include unanswerable questions, conflicting versions, difficult tables, cross-document questions and adversarial content.

12.3 Stage 2: controlled prototype

Implement source preservation, catalogue, permission-aware retrieval, typed candidate facts and the review workbench. Keep all outputs internal and clearly marked as candidates. Disable external communication and broad tool authority.

12.4 Stage 3: validation and security

Measure extraction, retrieval, conflict, abstention, calculation, citation and access controls. Run security and cross-matter tests. Remediate failures and repeat on a fresh locked set. Thresholds require management and risk-owner approval.

12.5 Stage 4: shadow operation

Run assisted and current workflows in parallel. Review every packet. Compare time, quality, correction, queue and reviewer measures. Do not use shadow output for external communications or final investment authority.

12.6 Stage 5: restricted production

Release defined low- or medium-consequence tasks to named users and sources. Maintain full review for material facts and all external or investment outputs. Monitor drift, exceptions and capacity. Preserve immediate manual fallback.

12.7 Stage 6: scale, redesign or retire

Scale one validated task and document class at a time. Revalidate after model, prompt, parser, policy or source-distribution changes. Redesign tasks that generate excessive exceptions. Retire a use case when risk, cost or reviewer burden exceeds approved value.

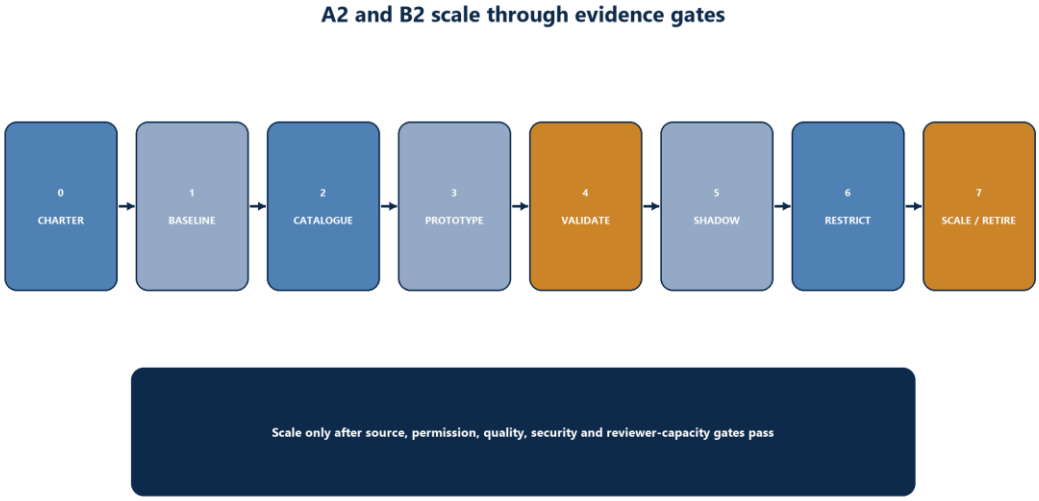

Figure 7. Gated adoption roadmap

Source: Matchpoint adoption framework.

Stage	Entry evidence	Exit evidence	Stop condition
0 Charter	Named decision and sponsor	Approved scope, roles and risk boundary	Authority or permitted purpose unclear
1 Baseline	Representative current cases	Observed baseline and locked corpus	Sample excludes material complexity
2 Prototype	Approved architecture and vendors	Candidate packets with full lineage	Source or access control fails
3 Validate	Answer key and test plan	Approved quality/security results	Critical miss, leakage or forced answer
4 Shadow	Review capacity and manual comparator	Observed quality and productivity evidence	Reviewer queue or correction rises materially

Stage	Entry evidence	Exit evidence	Stop condition
5 Restrict	Named tasks, users and rollback	Stable monitored operation	Drift, incident or unsupported release
6 Scale/retire	Approved operating evidence	Expanded scope or documented withdrawal	Risk-adjusted case is not approved

13. LIMITATIONS AND RESEARCH AGENDA

13.1 Evidence limitations

No Matchpoint client document, data room, investment paper, investor questionnaire, fund model, vendor configuration, model output, reviewer study or operating-cost record was supplied for T15. The paper therefore establishes a design and evaluation framework. It does not establish current Matchpoint capability, client performance, regulatory compliance, time saving, ROI or commercial impact.

The source set combines law, regulator material, standards, industry templates and technical research. Standards and technical papers describe general methods. Jurisdictional and professional applicability requires current case-specific review. Online sources can change after the evidence cut-off.

13.2 Technical limitations

Document layouts, languages, scans, spreadsheets, handwriting, encryption and custom financial structures create distribution shift. Benchmark performance does not transfer automatically. Retrieval and generation metrics can miss decision-specific harm. Human review can suffer inconsistency and automation bias. Vendor systems can change without transparent model details.

13.3 Economic limitations

The productivity table is an unverified illustrative management scenario. It excludes adoption friction, procurement, integration, security, validation, reviewer training, correction and vendor-switching costs except where illustrative values are explicitly shown. Capacity released may not convert to cost reduction or revenue. Attributed revenue, cost reduction and loss reduction remain USD 0.

13.4 Research agenda

A representative pilot should answer five questions. First, which question and document classes achieve the required critical-error and abstention performance? Second, how much reviewer time moves from search to judgement? Third, which access and prompt-injection tests expose architectural weakness? Fourth, how stable are results across funds, formats, languages and model versions? Fifth, what observed cost per accepted packet supports or rejects restricted production?

Future evaluation should publish the corpus design, question contracts, source authority, metrics, thresholds, model/parser versions, reviewer protocol and exclusions. Client confidentiality may prevent release of data; aggregated methodology and limitation reporting can still support governance.

14. CONCLUSION

AI document intelligence can turn a data room into a controlled evidence workflow when source identity, version, context, permissions, calculation, conflict, abstention and human authority are built into the operating model. The central deliverable is a quality-adjusted accepted decision-evidence packet. It connects a decision question to the exact sources, candidate facts, calculations, unresolved conflicts and named approval that support a released answer.

For A2 family-office CIOs, the framework supports manager-screening, diligence planning, track-record preparation and source-linked committee papers. For B2 GCC fund managers, it supports data-room readiness, controlled investor responses and substantiated communications. Each use remains bounded by governing documents, professional review, confidentiality and the approved decision process.

The practical sequence is charter, baseline, catalogue, prototype, validation, shadow operation, restricted release and evidence-led scale or retirement. Productivity should be measured per accepted packet at approved quality and security gates. Commercial attribution remains USD 0 until observed results and causal attribution are approved.

DECLARATIONS

Prepared by. Matchpoint Partners. Author attribution is pending CK approval.
Evidence cut-off. 1 August 2026.
Author-status note. No named-person authorship is asserted in this version.
Data availability. No client data room, fund document set, investor file, model output or pilot dataset was supplied or analysed. The paper relies on the cited public sources and explicitly labelled illustrative assumptions.
Commercial attribution. Attributed revenue: USD 0. Attributed cost reduction: USD 0. Attributed loss reduction: USD 0.
Competing interests. Matchpoint Partners may advise investment firms, fund managers and sponsors on capital, strategy, technology and transaction matters. No client outcome is represented in this paper.
Use boundary. General research for professional audiences. This paper is not investment, legal, regulatory, privacy, cybersecurity, accounting, tax, audit, valuation or technology advice; it is not an offer, solicitation, recommendation or promise of results.

REFERENCES

[1] Institutional Limited Partners Association. Due Diligence Questionnaire. Updated 2021. https://ilpa.org/resources-tools/resource-library/due-diligence-questionnaire/
[2] Institutional Limited Partners Association. ILPA Reporting Template v2.0: Suggested Guidance. 2025. https://ilpa.org/wp-content/uploads/2025/01/ILPA-Reporting-Template-v.-2.0-Suggested-Guidance.pdf
[3] Institutional Limited Partners Association. ILPA Performance Template: Suggested Guidance, Granular Methodology. 2025. https://ilpa.org/wp-content/uploads/2025/01/ILPA-Performance-Template-Suggested-Guidance-Granular-Methodology.pdf
[4] Institutional Limited Partners Association. ILPA Principles 3.0. 2019. https://ilpa.org/wp-content/uploads/2019/06/ILPA-Principles-3.0_2019.pdf
[5] Institutional Limited Partners Association. ILPA Portfolio Company Metrics Template. Accessed 1 August 2026. https://ilpa.org/industry-guidance/templates-standards-model-documents/ilpa-templates-hub/ilpa-portfolio-company-metrics-template/
[6] U.S. Securities and Exchange Commission. Investment Adviser Marketing: A Small Entity Compliance Guide. https://www.sec.gov/resources-small-businesses/small-business-compliance-guides/investment-adviser-marketing
[7] U.S. Securities and Exchange Commission, Division of Investment Management. Marketing Compliance Frequently Asked Questions. Updated 15 January 2026. https://www.sec.gov/rules-regulations/staff-guidance/division-investment-management-frequently-asked-questions/marketing-compliance-frequently-asked-questions
[8] U.S. Securities and Exchange Commission. Rule 204-2, Books and Records to Be Maintained by Investment Advisers: Reverted Rule Text. https://www.sec.gov/files/investment/pfa-vacatur-reverted-rule-text.pdf
[9] U.S. Securities and Exchange Commission. Electronic Recordkeeping by Investment Companies and Investment Advisers; Correction. 2001. https://www.sec.gov/rules-regulations/2001/05/electronic-recordkeeping-investment-companies-investment-advisers-correction
[10] U.S. Securities and Exchange Commission. Frequently Asked Questions on Form ADV and IARD. https://www.sec.gov/about/divisions-offices/division-investment-management/electronic-filing-investment-advisers-iard/frequently-asked-questions-form-adv-iard
[11] U.S. Securities and Exchange Commission. Inline XBRL. https://www.sec.gov/data-research/structured-data/inline-xbrl
[12] U.S. Securities and Exchange Commission. Standard Taxonomies: Investment Companies. Updated 16 March 2026. https://www.sec.gov/data-research/structured-data/taxonomies-schemas/standard-taxonomies/investment-companies
[13] Dubai Financial Services Authority. Collective Investment Funds. https://www.dfsa.ae/what-we-do/collective-investment-funds
[14] Abu Dhabi Global Market Financial Services Regulatory Authority. Regulations and Rules. https://www.adgm.com/operating-in-adgm/additional-obligations-of-financial-services-entities/supervision/regulations-and-rules/
[15] Abu Dhabi Global Market Financial Services Regulatory Authority. Getting Started Guide with FSRA. https://assets.adgm.com/download/assets/Getting%2BStarted%2BGuide%2Bwith%2BFSRA.pdf/a8c5b07a589a11efa1111ed8b674bcd6
[16] United Arab Emirates. Federal Decree-Law No. 45 of 2021 Concerning the Protection of Personal Data. https://uaelegislation.gov.ae/en/legislations/1972
[17] Abu Dhabi Global Market. Data Protection Regulations 2021. https://www.adgm.com/documents/office-of-data-protection/resources/adgm-data-protection-regulations-2021-updated.pdf
[18] European Union. Regulation (EU) 2016/679, General Data Protection Regulation. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679
[19] European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. https://eur-lex.europa.eu/eli/reg/2024/1689/
[20] International Organization for Standardization. ISO 15489-1:2016 Information and documentation; Records management; Part 1: Concepts and principles. https://www.iso.org/standard/62542.html
[21] International Organization for Standardization. ISO 18829:2017 Document management; Assessing ECM/EDRM implementations; Trustworthiness. https://www.iso.org/standard/63513.html
[22] International Organization for Standardization. ISO/IEC 27001:2022 Information security management systems. https://www.iso.org/standard/27001
[23] National Institute of Standards and Technology. SP 800-207 Zero Trust Architecture. 2020. https://doi.org/10.6028/NIST.SP.800-207
[24] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). 2023. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10
[25] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. 2024. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
[26] National Institute of Standards and Technology. Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, NIST AI 100-2e2025. 2025. https://doi.org/10.6028/NIST.AI.100-2e2025
[27] International Organization for Standardization. ISO/IEC 42001:2023 Information technology; Artificial intelligence; Management system. https://www.iso.org/standard/42001
[28] Xu, Y. et al. LayoutLM: Pre-training of Text and Layout for Document Image Understanding. 2020. https://www.microsoft.com/en-us/research/wp-content/uploads/2020/02/layoutlm.pdf
[29] Xu, Y. et al. LayoutLMv2: Multi-modal Pre-training for Visually-Rich Document Understanding. 2020. https://arxiv.org/abs/2012.14740
[30] Mathew, M., Karatzas, D. and Jawahar, C.V. DocVQA: A Dataset for VQA on Document Images. WACV 2021. https://openaccess.thecvf.com/content/WACV2021/html/Mathew_DocVQA_A_Dataset_for_VQA_on_Document_Images_WACV_2021_paper.html
[31] Tito, R. et al. Document Collection Visual Question Answering. 2021. https://arxiv.org/abs/2104.14336
[32] Wang, D. et al. DocLLM: A Layout-Aware Generative Language Model for Multimodal Document Understanding. ACL 2024. https://aclanthology.org/2024.acl-long.463.pdf
[33] Kim, G. et al. OCR-free Document Understanding Transformer. ECCV 2022. https://www.ecva.net/papers/eccv_2022/papers_ECCV/papers/136880493.pdf
[34] Rajpurkar, P., Jia, R. and Liang, P. Know What You Don't Know: Unanswerable Questions for SQuAD. 2018. https://arxiv.org/abs/1806.03822
[35] Noy, S. and Zhang, W. Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. Science 381, 187-192. 2023. https://doi.org/10.1126/science.adh2586
[36] Brynjolfsson, E., Li, D. and Raymond, L.R. Generative AI at Work. NBER Working Paper 31161. 2023. https://www.nber.org/papers/w31161
[37] Dell'Acqua, F. et al. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. Harvard Business School Working Paper 24-013. 2023. https://www.hbs.edu/ris/download.aspx?name=24-013.pdf
[38] IFRS Foundation. IFRS Accounting Taxonomy. https://www.ifrs.org/issued-standards/ifrs-taxonomy/
[39] XBRL International. Inline XBRL 1.1. Updated 14 July 2026. https://specifications.xbrl.org/work-product-index-inline-xbrl-inline-xbrl-1.1.html
[40] XBRL International. XBRL Specifications. https://www.xbrl.org/the-standard/what/specifications/

APPENDIX A. DECISION-EVIDENCE PACKET MINIMUM FIELDS

Field group	Minimum fields
Question contract	Question ID, text, entity, vehicle, period, currency, purpose, output schema and reviewer
Source identity	Document ID, family, version, hash, origin, received time, access class and status
Evidence span	Page/cell coordinates, extracted content, layout context and retrieval method
Candidate fact	Type, value, unit, period, entity, basis, confidence and extraction version
Calculation	Approved inputs, formula, code version, result and reconciliation
Conflict	Competing propositions, sources, class, owner and resolution
Review	Reviewer, decision, reason, time and override
Release	Recipient, purpose, approved wording, version, expiry and correction link

APPENDIX B. VALIDATION SCORECARD

Test family	Required reporting
Corpus	Document families, funds, languages, formats, periods, access classes and exclusions
Extraction	Field-level exact match; critical-field error; table and OCR strata
Retrieval	Required-span recall; review-set precision; version and access precision
Answering	Claim evidence, abstention, conflict recall and unsupported-claim rate
Calculation	Input match, unit/period checks, reproduction and reconciliation
Review	Acceptance, override, disagreement, time and queue age
Security	Cross-matter isolation, prompt injection, export and log-access tests
Operations	Latency, availability, correction, drift and rollback evidence

APPENDIX C. VENDOR AND ARCHITECTURE QUESTIONS

1. Which original and derivative data does the service store, where, for how long and under whose control?
2. Does any selected configuration permit provider training, human review or product improvement from customer content?
3. How are matter, vehicle, investor and purpose permissions enforced across source, index, cache, model context, output and logs?
4. Which parser, OCR, embedding and model versions processed each packet, and how can a prior result be reproduced?
5. How does the system preserve PDF regions, table structure, spreadsheet formulas and structured-data contexts?
6. How are superseded documents, investor-specific side letters and conflicting sources represented?
7. What happens when evidence is absent, inaccessible, ambiguous or outside professional authority?
8. Which tools can the model call, and which non-model controls prevent content-based instructions from changing system behaviour?
9. How are model, prompt and parser updates tested against a locked regression corpus and security harness?
10. How can the firm export originals, metadata, evidence packets, approvals and audit history in a usable format?
11. What subcontractors, subprocessors, transfer locations, encryption controls and incident processes apply?
12. How are deletion, legal hold, retention, correction and recipient notification implemented and evidenced?

APPENDIX D. RELEASE CHECKLIST

1. Confirm caller, recipient, matter, vehicle and permitted purpose.
2. Confirm every source is current for the question or explicitly labelled otherwise.
3. Resolve or disclose material conflicts and unavailable evidence.
4. Validate entity, unit, currency, period, gross/net and actual/forecast basis.
5. Reproduce every material calculation from approved typed inputs.
6. Open each material citation at the exact page, region or cell.
7. Confirm qualified review for legal, compliance, accounting, tax, valuation, cyber and investment matters.
8. Confirm qualifications, legends and marketing substantiation where applicable.
9. Record approval, version, release time, recipient, expiry and correction dependencies.
10. Preserve the released packet and maintain a practical correction route.

APPENDIX E. MANAGEMENT QUESTIONS

1. Which A2 or B2 decision should the first pilot improve, and who owns that decision?
2. Which document and question classes are in scope, and which are explicitly excluded?
3. Which data may enter the system, under which permissions, locations and retention rules?
4. Which source families and version-precedence rules are approved for each material question?
5. Which fields or claims are critical, and which error, conflict and abstention thresholds require approval?
6. Which reviewers have legal, compliance, finance, valuation, cyber and investment authority?
7. What current baseline will be observed for preparation, review, correction, queue and cost?
8. What vendor, integration, validation, security, training and operating costs are approved?
9. Which production tasks may use sampled review, and which require full review?
10. Which release channels remain disabled until explicit approval?
11. What incident, rollback, correction and recipient-notification process applies?
12. What evidence would cause management to scale, redesign or retire the use case?