

Building a Proprietary Knowledge Engine for an Advisory Firm

A governed decision-memory architecture for family-office investment teams and GCC fund managers

Prepared by Matchpoint Partners

Author attribution pending CK approval

August 2026

ABSTRACT

Background. Advisory firms create knowledge through research, client work, investment decisions and market relationships, yet fragmented files, messages and memories often prevent safe reuse. **Objective.** This paper develops a proprietary knowledge-engine architecture for A2 family-office CIOs and heads of alternatives and B2 GCC fund managers raising capital. **Approach.** The analysis reviews 40 primary, authoritative and clearly labelled sources available through 1 August 2026 across organisational knowledge, retrieval, generative AI, productivity, provenance, privacy, security and financial-services governance. **Findings.** A document repository becomes a decision engine only when sources, claims, permissions, decisions, outcomes and supersession are governed as one evidence chain. Retrieval quality, citation precision, faithfulness, freshness, conflict detection, abstention and permission leakage require measured hard gates. **Implications.** Firms should begin with a permissioned decision-memory pilot, independent evaluation and human approval. Experimental productivity evidence is task-specific and does not establish Matchpoint or client economics. All worked inputs are unverified illustrative management assumptions; attributed Matchpoint or client revenue, cash cost reduction and loss reduction remain USD 0 until approved observed evidence exists.

Highlights

Asset. The proprietary asset is governed decision memory with traceable evidence, ownership and outcomes.

Users. A2 and B2 teams require separate permissions, workflows and decision rights within a shared control model.

Architecture. Immutable originals, metadata, access control, hybrid retrieval, provenance, generation, review and audit remain separable.

Evaluation. Retrieval, citations, faithfulness, abstention, freshness, conflicts and permission isolation require independent tests.

Economics. Recognised value follows approved observed capacity, cost, revenue or loss evidence; generic productivity results remain external evidence.

JEL Classification: D83, G23, G24, L84, M15, O32, O33

Keywords: proprietary knowledge engine, advisory firm, organisational memory, retrieval-augmented generation, decision provenance, family office, fund manager, investment committee, due diligence, document intelligence, knowledge governance, model risk

Research paper only. This document is not investment, legal, regulatory, accounting, audit, tax, privacy, cybersecurity, employment, valuation or technology advice. All worked inputs, thresholds, times, rates and economics are unverified illustrative management assumptions.

1. INTRODUCTION

An advisory firm creates knowledge each time it researches a market, challenges an investment case, prepares an investment-committee memorandum, answers a limited partner, structures a mandate, records a meeting or observes an outcome. Much of that knowledge remains dispersed across email, shared drives, customer-relationship systems, data rooms, presentation files and individual memory. The commercial and fiduciary problem appears when a team cannot find the relevant evidence, cannot establish which statement is current, cannot see why a decision was made, or cannot reuse prior work without crossing a client or confidentiality boundary.

This paper addresses **Building a Proprietary Knowledge Engine for an Advisory Firm** for two Matchpoint Partners target audiences. **A2 Family-Office CIOs and Heads of Alternatives** are investment professionals at single-family offices, multi-family offices, private-wealth firms and external asset managers across the GCC and other major financial centres. **B2 GCC Fund Managers and GPs Raising Capital** are emerging and established private-equity, private-credit and venture-capital managers preparing institutional fundraising, diligence and investor-relations materials. Both audiences operate through evidence-intensive decisions. Their knowledge requirements overlap, while their permissions, approval paths and uses of information remain distinct.

The governing question is: **how can an advisory firm turn dispersed work product into reusable decision memory while preserving evidence, permissions, accountability and commercial truth?** A useful answer requires more than a search interface or a language model. The system must preserve original sources; represent claims, entities, decisions and outcomes; enforce access before retrieval; show citations; identify conflicts and stale material; support abstention; record human approval; and retain an audit trail.

The knowledge-based view of the firm treats knowledge integration as a central organisational capability [1]. Research on organisational knowledge creation, knowledge-management systems, transfer, absorptive capacity and internal stickiness explains why useful knowledge is distributed, partly tacit and difficult to move reliably [2-8]. Retrieval-augmented generation, dense retrieval, late interaction, long-context research, self-reflection, graph-based retrieval and evaluation frameworks add new technical tools [9-20]. These tools change the feasible interface to a firm's corpus. They do not remove the underlying requirements for ownership, provenance, confidentiality, validation or decision rights.

The paper defines a **proprietary knowledge engine** as a governed system that connects source evidence to claims, people, organisations, engagements, decisions, approvals and observed outcomes. Proprietary value arises from the firm's authorised evidence and operating history, its structured interpretation, and its controlled feedback loop. The model provider and retrieval software can be replaceable components. The durable asset is the permissioned decision record and the discipline that keeps it current.




Figure 1. The proprietary knowledge compounding loop

Source: Matchpoint evidence framework; organisational-knowledge and provenance context [1-8,27].

The contribution is a practical operating framework. It provides:

- a definition of proprietary knowledge and decision memory;
- an A2 and B2 decision perimeter;

- a minimum knowledge model and provenance record;
- a before-and-after workflow;
- a controlled reference architecture and tool-stack logic;
- retrieval, generation and evaluation gates;
- privacy, security and regulatory controls;
- a productivity and value-attribution bridge;
- unverified illustrative management scenarios;
- a ninety-day adoption roadmap; and
- a claims register for external communications.

The analysis reviews 40 primary, authoritative and clearly labelled sources available through 1 August 2026. Experimental productivity findings are reported within their source settings. Technical research is treated as evidence for specific components and evaluation methods. Regulatory sources apply according to jurisdiction and activity. No approved observed Matchpoint or client evidence was supplied for T24 revenue, cash cost reduction or loss reduction. Those attributed values therefore remain USD 0.

2. WHAT A PROPRIETARY KNOWLEDGE ENGINE IS

2.1 From repository to decision memory

A repository preserves files. A knowledge engine connects authorised evidence to a current decision. The distinction is operational.

Capability	File repository	Search layer	Proprietary knowledge engine
Stores originals	Yes	Usually	Yes, with integrity metadata
Finds text	Basic	Stronger	Permissioned and decision-specific
Represents claims	Rarely	Indirectly	Explicit claim records
Shows provenance	File-level	Variable	Source, passage, owner and date
Identifies supersession	Manual	Limited	Explicit current and superseded states
Links decisions and outcomes	Rarely	Rarely	Core knowledge object
Enforces approval	Folder-based	Variable	Workflow and decision-level
Learns from outcomes	No	No	Controlled feedback process

Grant describes firms as institutions for integrating specialist knowledge [1]. Nonaka explains organisational knowledge creation through interaction between tacit and explicit knowledge [2]. Alavi and Leidner organise knowledge-management systems around creation, storage and retrieval, transfer and application [3]. Argote and Ingram show that knowledge transfer can occur through movement of people, tools and networks [4]. These foundations support a design that captures context and use, rather than treating documents as interchangeable text.

The engine should preserve six connected layers:

1. **Source layer.** The immutable or versioned original, including ownership, licence, confidentiality, jurisdiction and retention.
2. **Claim layer.** A bounded statement linked to supporting or contradicting passages, date, confidence and reviewer.
3. **Entity layer.** People, organisations, funds, assets, sectors, jurisdictions, services and engagements.
4. **Decision layer.** The question, alternatives, evidence, assumptions, owner, approval and date.
5. **Outcome layer.** The later observation that can confirm, qualify or contradict the decision rationale.
6. **Permission layer.** The actor, purpose, client wall, engagement boundary and permitted action.

2.2 What proprietary means

The label **proprietary** requires a documented basis. A public market report remains public evidence even when it is indexed internally. A model-generated summary remains derived content. A firm's authorised interview notes, diligence findings, client-approved work product, structured decision records, relationship history and outcome observations can form proprietary knowledge when the firm has the right to retain and use them for the stated purpose.

Knowledge class	Example	Default treatment
Public authoritative	Regulator publication	Retain citation, version and access date
Licensed third-party	Research subscription	Enforce licence and redistribution terms
Client-confidential	Data-room or mandate material	Client wall and purpose limitation
Internal operating	Approved template or playbook	Role-based access and version owner
Relationship knowledge	Meeting note or contact preference	Lawful basis, purpose and quality review
Generated derivative	Draft summary or answer	Label, cite, review and expire as appropriate
Decision memory	Approved rationale and outcome	Preserve authority, date and supersession

The system must avoid claiming ownership over public or client-owned information. Its proprietary advantage can reside in structure, verified interpretation, relationship context, workflow and accumulated decision outcomes. Rights metadata belongs in the minimum record.

2.3 Knowledge creation and transfer

Cohen and Levinthal define absorptive capacity as the ability to recognise, assimilate and apply external knowledge [5]. Szulanski identifies internal stickiness in knowledge transfer [6]. March distinguishes exploration of new possibilities from exploitation of existing knowledge [7]. Kogut and Zander link combinative capability to the creation and transfer of knowledge [8]. These concepts have direct system implications.

The engine should support exploration through new research, dissent and weak-signal capture. It should support exploitation through approved templates, prior answers and repeatable evidence packs. A single ranked answer can suppress useful disagreement. The interface should expose supporting and contradicting evidence, older decisions, superseding events and the identity of the responsible reviewer.

2.4 Evidence classes

Every output should carry an evidence state:

State	Meaning	Permitted use
Source	Original evidence preserved	Retrieval and review
Extracted	Machine- or human-extracted content	Search with source link
Generated	Model-created draft	Review only
Reviewed	Human checked for stated use	Controlled reuse
Approved	Authorised for a defined purpose	Use within purpose and validity period
Superseded	Replaced by newer approved evidence	Historical context only
Disputed	Conflicting evidence unresolved	Escalate; no definitive claim

The system should abstain when the evidence state does not support the requested action. A well-written answer does not change the evidence class.

3. THE A2 AND B2 DECISION PROBLEM

3.1 A2 family-office investment teams

A2 teams evaluate funds, direct investments, co-investments, managers, service providers and strategic portfolio questions. Their information can include family identity, legal entities, liquidity needs, portfolio exposures, manager materials, tax and jurisdictional context, investment-committee papers and private relationship history. The system must preserve a family or client perimeter before it offers convenience.

Common A2 knowledge-engine questions include:

- Which evidence supported the previous manager decision?
- Which terms, risks and open issues changed since the last committee?
- Where has this sponsor, executive or adviser appeared before?
- Which portfolio exposures interact with the proposed allocation?

- Which assumptions remain unverified?
- Which paper or person owns the current answer?
- What information can be reused across family entities or external advisers?

The engine can assist with research assembly, comparison and draft preparation. The investment committee or authorised delegate retains decision authority. Suitability, allocation, legal, tax and regulatory conclusions require the relevant accountable professionals.

3.2 B2 GCC fund managers raising capital

B2 teams repeatedly answer limited-partner questions about strategy, team, track record, attribution, valuation, risk, operations, ESG, governance, service providers, portfolio construction and pipeline. Answers often exist across a private-placement memorandum, due-diligence questionnaire, data room, finance records, portfolio-company files and prior correspondence. Inconsistent or stale answers can damage credibility and create regulatory or contractual exposure.

Common B2 questions include:

- What is the current approved response to this DDQ question?
- Which source supports each track-record or operating claim?
- Which answer can be reused for this investor and jurisdiction?
- What changed after the latest fund close, valuation or team event?
- Which materials contain hypothetical, gross, net or target performance?
- Who approved the external wording and when does it expire?
- Which investor interaction or objection should inform the next response?

The engine can accelerate evidence assembly and identify prior approved language. Fund counsel, compliance, finance and the authorised manager remain responsible for external disclosure.

3.3 Shared model, separate rights




Figure 2. A2 and B2 decision perimeter

Source: Matchpoint ICP and decision-rights framework.

Object	A2 primary use	B2 primary use	Shared control
Entity	Family, manager, asset	LP, fund, portfolio company	Identity resolution and access
Claim	Investment thesis or risk	Fundraising or track-record statement	Evidence and approval
Decision	Allocation or diligence	Disclosure or fundraising action	Owner, date and rationale
Outcome	Performance, event, lesson	LP response, close, exception	Observation and attribution
Relationship	Sponsor and adviser history	LP and intermediary history	Purpose and privacy

Object	A2 primary use	B2 primary use	Shared control
Template	IC memorandum	DDQ or data-room index	Version and reuse boundary

The A2 and B2 systems may share architecture, taxonomy and evaluation methods. They require separate client walls, access lists, purposes, audit trails and external-publication controls. Cross-engagement retrieval should be denied by default until rights and purpose are established.

3.4 Decision-rights matrix

Action	Knowledge owner	System owner	Reviewer	Decision authority
Add source	Confirms rights and context	Ingests and records	Samples quality	No decision effect
Approve claim	Supplies evidence	Preserves lineage	Challenges support	Authorised business owner
Answer research question	Frames purpose	Retrieves and drafts	Checks evidence	Named professional
Reuse client material	Confirms permission	Enforces wall	Privacy/legal as required	Engagement owner
Publish externally	Supplies approved content	Records version	Compliance/legal as required	Authorised signatory
Change model or retrieval	Defines risk	Implements	Independent validation	Change authority

System access does not confer investment, disclosure or publication authority.

4. THE MINIMUM KNOWLEDGE MODEL

4.1 Source record

The source record is the integrity anchor. It should contain:

Field	Purpose
Source ID and checksum	Detect replacement or corruption
Original location and owner	Establish custody
Created, received and effective dates	Support the decision-time view
Confidentiality and client	Enforce wall
Licence or use right	Restrict copying and distribution
Jurisdiction and personal-data class	Route privacy control
Retention and deletion rule	Manage lifecycle
Parser and version	Reproduce extraction
Superseding source	Identify current record
Review status	Bound permitted use

Immutable does not require permanent retention. It means the retained evidentiary object cannot be silently changed. Deletion and legal-hold rules remain governed.

4.2 Claim record

A claim is the smallest unit that can be supported, contradicted, approved or superseded. A useful record includes exact wording, source passages, source date, subject entities, numerical unit, calculation method, conditions, evidence status, reviewer, validity period and external-use permission.

Claims require granularity. The sentence "Fund performance was strong" is too vague. A usable record states the metric, period, gross or net status, currency, benchmark, valuation date, calculation method and source. Marketing and fundraising claims should also identify the intended audience and required disclosures.

4.3 Decision and outcome record

Decision memory should answer four questions:

1. What decision was made?

2. What evidence and assumptions were available at the time?
3. Who challenged and approved it?
4. What happened later?

Decision field	Required content
Question	Decision that required authority
Alternatives	Options considered, including stop or defer
Evidence	Supporting and contradicting claims
Assumptions	Explicit, dated and owned
Dissent	Material challenge or unresolved issue
Authority	Approver and governance basis
Date and horizon	Decision time and expected review time
Outcome	Later observation and evidence
Lesson	Approved update to future practice
Supersession	Later decision that changes current state

Outcome capture closes the compounding loop. A system that stores decisions without later observations can reproduce prior reasoning while missing the information needed to improve it.

4.4 Provenance model

The W3C PROV data model distinguishes entities, activities and agents and describes how provenance relationships can be represented [27]. An advisory implementation can translate that model into sources, transformations, claims, drafts, reviews, approvals and decisions.

Each generated answer should expose:

- the source ID and passage;
- the retrieval time and query purpose;
- the transformation or extraction version;
- the model and prompt or workflow version where relevant;
- the permissions applied;
- the reviewer and approval state; and
- any later correction or supersession.

Provenance supports investigation and challenge. It does not establish truth by itself. A perfectly traceable source can still be wrong, outdated or inapplicable.

5. BEFORE AND AFTER: THE OPERATING WORKFLOW

5.1 Fragmented baseline

The baseline process often depends on a senior person's memory of where information lives. A request triggers email searches, folder browsing, colleague messages, copied answers and manual reconciliation. The team may find several versions and select the most plausible one. Work is repeated because the prior rationale was not captured or cannot be reused under the current permission.

Baseline failure	Observable symptom	Risk
Fragmentation	Same question answered from multiple files	Inconsistency
Tacit dependence	Work stops when one person is unavailable	Key-person exposure
Weak lineage	Numbers copied without calculation source	Unsubstantiated claim
Staleness	Old deck used after fund or portfolio change	Misstatement
Permission ambiguity	Client material appears in broad search	Confidentiality breach
Outcome loss	Prior decision rationale never revisited	Repeated error

Baseline failure	Observable symptom	Risk
Duplicate effort	Analysts recreate prior work	Capacity loss

5.2 Governed target workflow

Before and after: advisory knowledge workflow




Figure 3. Before and after: advisory knowledge workflow

Source: Matchpoint operating-model framework.

The target workflow begins with a decision question and authorised purpose. The system filters the corpus by actor and purpose before retrieval. It returns source passages and structured claims, identifies conflicts and freshness, drafts a response with citations, and routes the result to the required reviewer. Approval produces a versioned reusable answer for a defined purpose and period. The decision and later outcome become new governed records.

Step	System action	Human control	Evidence retained
Frame	Record question, user and purpose	Confirm decision boundary	Query charter
Filter	Apply client, role and purpose permissions	Approve exceptional access	Policy decision
Retrieve	Search, rerank and expand relevant evidence	Challenge omissions	Ranked source set
Draft	Generate bounded synthesis with citations	Edit or reject	Draft and model record
Validate	Test support, conflicts, numbers and freshness	Professional review	Exceptions and resolution
Approve	Freeze permitted version	Authorised sign-off	Approval and expiry
Reuse	Serve within the approved context	Confirm current applicability	Usage log
Learn	Add outcome and supersession	Approve lesson	Updated decision memory

5.3 Designed abstention

The workflow should return "insufficient approved evidence" when the corpus, permission or freshness gate fails. Abstention is an intended control state. Escalation paths should identify the missing owner, evidence or approval.

Examples include:

- current fund performance is unavailable at the approved reporting date;
- a client wall excludes the only relevant document;
- two approved sources conflict and no owner has resolved them;
- a personal-data purpose does not cover the requested use;
- a numerical claim lacks an approved calculation; or

- the external-use approval expired after a material event.

6. CONTROLLED ARCHITECTURE AND TOOL STACK

6.1 Architectural layers




Figure 4. Controlled knowledge-engine architecture

Source: Matchpoint architecture; retrieval, provenance, security and governance context [9-20,25-35].

The architecture separates six layers so that evidence and controls remain inspectable.

Source connectors

Connectors can ingest approved folders, email, CRM records, data-room exports, meeting records and authoritative external sources. Each connector requires a named owner, scope, refresh schedule and rights basis. Bulk connection of every available system expands exposure before use cases and permissions are defined.

Immutable originals and extraction

The system retains the authorised original or a durable reference, checksum, date and parser version. Optical character recognition and document extraction produce derivative text. Tables, footnotes, page structure and attachments should remain traceable because advisory claims frequently depend on layout and qualifications.

Metadata, access control and knowledge graph

Metadata includes entities, source class, client, engagement, jurisdiction, confidentiality, effective date, review status and retention. Access policy should operate before retrieval and after generation. A graph can connect entities, claims, sources, decisions and outcomes; it should preserve edge provenance and uncertainty.

Hybrid retrieval and reranking

Sparse lexical retrieval identifies exact names, phrases and codes. Dense retrieval identifies semantically similar passages [10]. Late-interaction methods such as ColBERT preserve token-level matching while supporting precomputed document representations [20]. A hybrid design can retrieve candidates, apply metadata and permissions, rerank, and expand through graph relationships.

Grounded assistance

Retrieval-augmented generation conditions a language model on selected evidence [9]. REALM, RETRO and Atlas illustrate different ways retrieval can support language-model training or inference [17-19]. The advisory engine should provide citations, delimit evidence, identify unsupported passages, and permit abstention. The model's pretrained knowledge should not silently substitute for the approved corpus in high-stakes answers.

Review, approval and audit

The workflow routes each output according to use. Internal research notes, investment-committee papers, client deliverables and public fundraising statements require different reviewers and retention. The system records the exact version that was approved and the evidence visible at that time.

6.2 Build, buy and portability

Component	Common sourcing choice	Proprietary boundary
Storage and search	Managed platform or cloud service	Source rights and metadata
Embeddings and reranker	Commercial or open model	Evaluation set and configuration
Language model	Commercial API, hosted or local	Approved workflow and evidence
Graph and taxonomy	Database plus firm design	Entity, claim and decision model
Access control	Identity provider and policy engine	Client and purpose rules
Evaluation	Framework plus human set	Gold questions and failure cases
Workflow	Case-management or custom layer	Review and authority design

Portability should cover original sources, metadata, embeddings where licensed, graph relationships, prompts, evaluation sets, approvals and logs. Vendor exit does not need to preserve every implementation detail. It should preserve the firm's evidence and ability to reconstruct material decisions.

6.3 Versioning and change control

A model, embedding, parser, prompt, retrieval parameter, taxonomy or access-policy change can alter answers. Each material change requires a defined risk class, regression suite, approver and rollback plan. NIST's Secure Software Development Framework supports integrating secure practices throughout development [28]. The AI Risk Management Framework and its generative-AI profile provide governance, mapping, measurement and management concepts for AI systems [25,26].

Change	Minimum test
Parser or OCR	Extraction accuracy on representative files
Embedding model	Retrieval recall and permission regression
Reranker	Ranking quality and latency
Language model	Faithfulness, citation and abstention
Prompt or workflow	Full end-to-end golden set
Taxonomy	Entity and filter consistency
Access policy	Adversarial cross-client isolation
Data source	Rights, quality, freshness and incident review

7. RETRIEVAL AND GENERATION: EVIDENCE AND LIMITS

7.1 Retrieval quality is corpus-specific

Dense Passage Retrieval reports improved open-domain question-answer retrieval against a strong BM25 system in its benchmark setting [10]. BEIR shows that zero-shot retrieval performance varies materially across heterogeneous datasets and tasks [11]. These findings support direct evaluation on the firm's own questions, sources and access rules. A vendor benchmark does not establish performance on private-placement memoranda, committee minutes, financial tables or multilingual GCC correspondence.

A minimum evaluation set should represent:

- exact entity and fund-name queries;
- synonyms, acronyms and transliteration;
- date-sensitive and supersession questions;
- numerical and table-based evidence;
- supporting and contradicting sources;
- permission-denied questions;
- questions with no approved answer;
- multi-document synthesis; and

- A2 and B2 workflows separately.

7.2 Context length and evidence placement

Lost in the Middle finds that model performance can degrade when relevant information is placed in the middle of long contexts in the tested settings [12]. A large context window therefore requires evidence selection and ordering. It does not remove retrieval design. Long data-room documents should be segmented with page, table, section and attachment lineage, then assembled according to the question.

7.3 Self-reflection and graph retrieval

Self-RAG combines retrieval and generation with reflection tokens in the reported research setting [13]. GraphRAG uses a graph-based index and community summaries to support questions over narrative private data in its reported approach [16]. These methods offer useful design hypotheses for advisory work: retrieve selectively, expose relationships and evaluate global as well as local questions. Their reported results do not certify a private advisory implementation.

Graph edges require evidence. A relationship between a person and company can be historic, indirect or disputed. A model-inferred edge should remain labelled and separate from an approved factual relationship.

7.4 Evaluation frameworks

RAGAS proposes reference-free evaluation signals for retrieval-augmented generation [14]. ARES evaluates context relevance, answer faithfulness and answer relevance with model judges and a smaller human-labelled set [15]. These frameworks can reduce the cost of repeated testing. High-stakes advisory evaluation still requires human ground truth, adversarial cases and direct inspection of failures.

Automated evaluators can share biases with the model under review. The control design should record evaluator version, human sample, disagreement rate and override process.

7.5 Citation discipline

A citation is useful when it supports the exact adjacent statement. Citation presence alone is insufficient. The engine should test:

- whether the cited source contains the claim;
- whether qualifications and dates are preserved;
- whether a numerical value and unit match;
- whether the source was approved for the use;
- whether another current source contradicts it; and
- whether the answer introduces material content without evidence.

The system should provide source passages for reviewer inspection. It should preserve page or cell references where the source format supports them.

8. PRIVACY, SECURITY AND FINANCIAL-SERVICES CONTROLS

8.1 Data protection and purpose

The UAE Government's official data-protection overview describes the federal personal-data framework and sector or free-zone regimes [31]. DIFC Data Protection Law No. 5 of 2020 and its amendments govern relevant processing in the DIFC [32]. ADGM Data Protection Regulations 2021 and official guidance apply to relevant ADGM processing [33]. Applicable obligations depend on the entity, data subject, purpose, location and activity. Qualified legal and privacy advice is required for implementation.

The engine should record:

- controller and processor roles;
- purpose and lawful basis as advised;
- categories of personal and special-category data;
- data-subject and jurisdiction scope;
- transfer mechanism where applicable;
- retention and deletion rules;
- processor and subprocessor inventory;
- access and disclosure logs; and

- impact assessment and incident route.

Family-office data can reveal wealth, family relationships, identity, residence, health, legal structures and investment activity. Fund-manager data can reveal LP identities, beneficial ownership, employees and portfolio-company information. Purpose limitation should be enforced at query time and output time.

8.2 Security threats

OWASP's Top 10 for LLM Applications describes risks including prompt injection, sensitive information disclosure, supply-chain exposure and insecure output handling [29]. MITRE ATLAS catalogues adversary tactics and techniques for AI-enabled systems [30]. NIST's generative-AI profile identifies risks and actions across the AI lifecycle [26]. These sources support a threat model that covers the corpus, retrieval layer, model, tools, users and outputs.

Threat	Advisory example	Control evidence
Prompt injection	Malicious text in a data-room document	Content isolation and instruction hierarchy tests
Poisoning	False relationship note added to corpus	Source authority, review and anomaly detection
Permission leakage	A2 client asks about another family	Pre-retrieval filter and adversarial isolation test
Sensitive disclosure	Generated answer includes personal data	Output policy, redaction and review
Insecure tool use	Draft triggers unauthorised CRM or email action	Least privilege and explicit approval
Model or vendor change	Output shifts after silent upgrade	Version pinning and regression gate
Citation fabrication	Link or passage does not support statement	Citation verifier and human sample
Logging exposure	Prompts retain client secrets	Controlled telemetry and retention

8.3 Financial-services governance and communications

The DFSA reported in its 2025 survey that 52 percent of 661 responding DIFC firms used AI; 60 percent reported some AI governance and 21 percent lacked clear accountability, within the survey definitions [34]. These figures describe the survey population and should not be generalised to every advisory firm. The DFSA's 4 June 2026 supervisory letter is an official current source for its regulatory expectations on AI risk management in the DIFC [35].

The SEC investment-adviser marketing rule requires, among other matters, substantiation and fair and balanced treatment of material risks and limitations for firms within its scope [36]. The SEC's 2024 risk alert reports examination observations on untrue or unsubstantiated statements, omissions, misleading inference and unfair presentation of risks or performance [37]. The SEC charged two investment advisers in 2024 over false and misleading statements about their purported use of AI [38]. IOSCO's 2025 consultation report describes AI use cases, risks and challenges in capital markets [39]. These sources support controlled claims and evidence retention.

A B2 manager should never allow a generated answer to convert an internal estimate into an approved external fact. The output should preserve gross or net status, time period, currency, benchmark, calculation method, hypothetical or target status, limitations and required disclosures. An A2 adviser should preserve the distinction between research assistance and authorised investment advice or decision.


8.4 Privacy risk management

NIST's Privacy Framework provides a voluntary enterprise risk-management structure for privacy [40]. A knowledge engine can map its processing to identify, govern, control, communicate and protect functions. The firm should maintain a data map and a record of material processing changes. Privacy testing should include subject access, correction, deletion, retention, transfer and incident workflows where applicable.

9. EVALUATION SCORECARD AND HARD GATES

9.1 Unit of evaluation

Evaluation starts with a real decision task, not a generic chatbot question. Each test case should specify user, purpose, permitted corpus, required answer, supporting evidence, contradictions, abstention condition and severity of failure.

**Figure 5. Knowledge-engine evaluation scorecard**

Source: Matchpoint evaluation framework; retrieval evaluation context [11-15].

Dimension	Example measure	Hard-fail condition
Retrieval relevance	Recall at k on approved evidence	Required source consistently absent
Citation precision	Supported cited claims / cited claims	Fabricated or materially wrong citation
Faithfulness	Material statements supported by context	Unsupported material statement
Freshness	Current source selected over superseded source	Superseded claim presented as current
Conflict detection	Known contradictions surfaced	Material conflict concealed
Abstention	Correct abstentions / required abstentions	Definitive answer without sufficient evidence
Permission isolation	Denied evidence retrieved or emitted	Any cross-client or prohibited disclosure
Numerical integrity	Values, units and periods reproduced	Material number or unit error
Reviewer acceptance	Approved without material correction	Acceptance below approved threshold
Latency and cost	End-to-end observed distribution	Breach of approved operating envelope

9.2 Thresholds

Thresholds should be approved for each use case after a baseline is measured. The paper does not supply universal numerical thresholds. A client-confidential retrieval test can require zero observed leakage in a defined adversarial suite, while broader relevance metrics can use approved tolerances. A zero observed failure in a finite sample does not prove impossibility of future failure.

9.3 Evaluation set governance

Control	Requirement
Ownership	Business owner and independent validator
Coverage	A2, B2, source types, languages and risk classes
Version	Frozen set plus controlled additions
Independence	Test cases not used solely for tuning
Adversarial set	Injection, poisoning, leakage and ambiguity
Human ground truth	Named reviewers and disagreement process
Regression	Run after material component change

Control	Requirement
Incident feed	Production failures become future tests

Evaluation data can itself contain confidential information. Access, retention and separation from vendor training require control.

9.4 Release decision

Average quality should not override a severe control failure. The release gate should block deployment when a material permission leak, unsupported external claim, fabricated citation, unresolved security vulnerability or missing accountable owner is observed. Exceptions require a named authority, rationale, compensating control and expiry.

10. OPERATING GOVERNANCE

10.1 Three control lines

A practical model separates delivery, independent challenge and oversight.

Line	Responsibilities
Use-case and system owners	Corpus rights, workflow, output review, monitoring and incident response
Independent risk, compliance, privacy, security or validation	Challenge, testing, policy and exceptions according to scope
Internal audit or equivalent oversight	Periodic assurance over governance and evidence

Smaller firms may assign several roles to fewer people. The responsibilities and conflicts should remain explicit.

10.2 Knowledge stewardship

Each high-value domain requires a steward responsible for taxonomy, source quality, current approved claims, review dates and unresolved conflicts. Technology teams cannot determine the business truth of an investment or fundraising statement. Business owners cannot approve access-control implementation without technical evidence. Joint ownership is required.

10.3 Human review

Meaningful review requires sufficient time, source access, competence and authority to reject the output. A reviewer who sees only the generated answer cannot evaluate provenance. A reviewer who is measured solely on throughput may be discouraged from challenging the system. Review metrics should include material corrections, abstentions, escalations and later incidents.

10.4 Incident and correction process

An incident process should cover confidentiality, privacy, security, factual error, misleading communication, incorrect permission, unsupported recommendation and loss of evidence. It should identify affected outputs and users, preserve logs, contain access, correct the record, assess notification duties, and add the failure to the regression suite.

Incident record	Minimum field
Detection	Time, reporter and channel
Scope	Sources, outputs, clients and users affected
Severity	Approved classification
Containment	Access or workflow action
Correction	Superseding content and recipients
Notification	Legal, regulatory, client or internal route
Root cause	Data, model, prompt, access, process or human
Prevention	Test, control and owner

11. UNVERIFIED ILLUSTRATIVE MANAGEMENT SCENARIOS

All inputs in this section are **unverified illustrative management assumptions**. They are provided to demonstrate evaluation logic. They do not describe Matchpoint or a client.

11.1 Case A: A2 investment-committee memory

An A2 team pilots the engine on a restricted set of historical manager-diligence materials. The use case retrieves prior evidence and produces a cited comparison for reviewer approval.

Illustrative input	Value	Status
Historical manager files	1,200	Unverified scenario
Gold evaluation questions	120	Unverified scenario
Authorised pilot users	8	Unverified scenario
Pilot period	12 weeks	Unverified scenario
Investment authority	None	Control assumption

The pilot measures retrieval recall, citation precision, permission isolation, material corrections and reviewer time. It does not make allocation decisions. A decision paper records whether the system recovered relevant historical challenges and whether reviewers accepted the evidence chain.

11.2 Case B: B2 DDQ answer library

A B2 manager creates a claim-level library from approved fund documents and historical DDQ responses. Every external answer requires current evidence and designated approval.

Illustrative input	Value	Status
Approved source documents	350	Unverified scenario
Historical DDQ questions	900	Unverified scenario
Reusable approved claims	480	Unverified scenario
External approvers	3	Unverified scenario
Automatic external sending	Disabled	Control assumption

The system flags track-record questions for finance verification, legal terms for counsel review, and expired answers after a material event. The pilot reports answer preparation time, material corrections, citation support and stale-answer detection. It does not publish or send answers automatically.

11.3 Case C: cross-client isolation

The firm creates synthetic client identities and adversarial questions designed to retrieve prohibited material.

Illustrative test	Cases	Required result
Direct other-client query	50	Deny and reveal no evidence
Indirect entity hint	50	Deny and reveal no evidence
Prompt injection in document	40	Ignore embedded instruction
Conflicting source	30	Surface conflict
No approved evidence	30	Abstain

Observed pilot results would require independent recording. No results are supplied in this paper.

12. PRODUCTIVITY, ECONOMICS AND ATTRIBUTION

12.1 External productivity evidence

Noy and Zhang report an experiment in professional writing tasks in which access to generative AI changed time and output quality in their study setting [21]. Brynjolfsson, Li and Raymond study 5,172 customer-support agents and report a 15 percent average increase in issues resolved per hour with heterogeneous effects across workers and tasks [22].

Dell'Acqua and co-authors report gains for tasks inside the tested model's capability frontier and quality reductions for a task outside it in their knowledge-worker experiment [23]. Wiles and co-authors describe generative AI as an exoskeleton in a 2024 working paper setting [24].

These findings support task-specific pilots and heterogeneous measurement. They do not establish the effect of a proprietary advisory knowledge engine. Advisory work differs in confidentiality, source structure, professional accountability, task frequency and error cost.

12.2 Evidence-gated value bridge




Figure 6. Evidence-gated productivity and ROI bridge

Source: Matchpoint business-case framework; external productivity context [21-24].

The business case should progress through five states:

1. **Task.** The use case and accepted output are defined.
2. **Time.** Baseline and pilot capacity are observed under comparable conditions.
3. **Quality.** Material corrections, error, reviewer acceptance and outcomes are measured.
4. **Realisation.** Management and finance approve how released capacity or improved outcome is used.
5. **Attribution.** Approved evidence connects the intervention to revenue, cash saving, loss reduction or other value.

12.3 Illustrative formulas

For capacity:

Illustrative gross capacity value = accepted hours released x approved loaded cost per hour.

For realised operating value:

Illustrative realised value = gross capacity value x approved realisation factor - incremental cash costs.

For revenue:

Illustrative attributed revenue = approved collected revenue causally linked to the intervention under the firm's attribution policy.

Each input requires observed evidence and approval. Time released can be redeployed without becoming a cash saving. Faster drafting can create rework or risk when quality declines. Pipeline or proposal value does not equal collected revenue.

12.4 Current attribution status

No approved observed Matchpoint or client financial evidence was supplied for T24.

Category	Attributed value
Matchpoint revenue	USD 0
Client revenue	USD 0
Matchpoint cash cost reduction	USD 0
Client cash cost reduction	USD 0
Loss reduction	USD 0

These values are evidence boundaries. They are not forecasts.

13. NINETY-DAY ADOPTION ROADMAP

13.1 Days 0-15: decision and authority

- Name one A2 or B2 decision workflow.
- Identify source, knowledge, system and decision owners.
- Define permitted users, clients, purposes and outputs.
- Record prohibited actions and publication boundaries.
- Establish legal, privacy, security and regulatory review routes.

Exit gate: signed use-case charter and no autonomous investment, disclosure or external-sending authority.

13.2 Days 16-30: corpus and rights

- Select a bounded source corpus.
- Record ownership, licence, client, confidentiality, jurisdiction and retention.
- Preserve originals and parser lineage.
- Build the entity, claim, decision and outcome taxonomy.
- Prepare synthetic and historical evaluation questions.

Exit gate: source-rights and data-quality review passes.

13.3 Days 31-45: retrieval baseline

- Establish lexical, dense and hybrid retrieval lanes.
- Apply permission filters before retrieval.
- Measure relevant-source recall and stale-source selection.
- Add reranking and graph expansion only where the baseline supports them.
- Record latency and cost.

Exit gate: the approved corpus can be searched without permission exceptions in the defined test.

13.4 Days 46-60: grounded assistance and evaluation

- Add cited generation for the bounded workflow.
- Test faithfulness, citations, conflicts, numbers and abstention.
- Run prompt-injection, poisoning and cross-client tests.
- Compare model and workflow variants on the same frozen set.
- Preserve all material failures.

Exit gate: no unresolved hard-fail condition.

13.5 Days 61-75: restricted pilot

- Train authorised users and reviewers.
- Operate with human approval and no autonomous external action.
- Measure task time, material correction, rejection, escalation and incidents.
- Interview users about missing context and workflow friction.
- Validate logs and correction paths.

Exit gate: independent pilot report and owner acceptance.

13.6 Days 76-90: committee gate

- Present favourable and unfavourable evidence.
- Decide whether to stop, redesign, extend or move to controlled production.
- Approve component versions, monitoring, incident and change processes.
- Approve productivity and financial attribution methods.
- Approve external claims separately.

Ninety-day knowledge-engine adoption roadmap



Figure 7. Ninety-day knowledge-engine adoption roadmap

Source: Matchpoint implementation framework.

The ninety-day outcome is a governed decision and evidence pack. It is not a promise of production deployment or financial return.

14. CLAIMS REGISTER

Claim	Evidence status	Permitted conclusion
Firms integrate distributed specialist knowledge	Supported by organisational research [1-8]	Knowledge integration is an operating capability
Retrieval can improve access to relevant evidence	Supported in bounded technical studies [9-20]	Evaluate on the firm's corpus
Long context removes the need for retrieval	Not established; evidence shows position effects [12]	Preserve retrieval and ordering tests
Graph retrieval can support broad corpus questions	Supported in reported GraphRAG approach [16]	Pilot with evidence-bearing edges
Automated RAG evaluation can reduce evaluation effort	Supported in bounded frameworks [14,15]	Retain human ground truth and challenge
Generative AI can improve productivity in some tasks	Supported in specific experiments [21-24]	Measure the advisory workflow directly
The engine can safely cross client boundaries	Not established	Prohibited without specific right and purpose
A generated answer is an approved fact	Not established	Require evidence state and approval
T24 proves Matchpoint or client financial value	Not established	Attributed value remains USD 0
T24 authorises automated investment or external disclosure	Not established	No authority conferred

15. LIMITATIONS AND CONCLUSION

15.1 Limitations

The paper provides an operating and evaluation framework. It does not implement or empirically test a Matchpoint knowledge engine. No private corpus, retrieval system, model, access-control policy or production workflow was evaluated for T24. Technical research uses different corpora, tasks, models and metrics, limiting direct transfer to advisory work. Organisational research describes general mechanisms that require firm-specific implementation.

The privacy and regulatory sources apply according to entity, jurisdiction, activity and facts [31-40]. The paper is not legal or regulatory advice. Applicable advisers must obtain qualified advice. Security threats evolve, and a finite evaluation cannot prove the absence of future failure.

The illustrative scenarios contain unverified management assumptions. The paper does not estimate adoption cost, staff capacity, model price, migration effort, revenue or savings. Named-person author attribution remains pending CK approval.

15.2 Conclusion

A proprietary knowledge engine can become a durable advisory asset when it preserves the evidence, rights, interpretation, decision and outcome chain. The useful unit is decision memory. Documents, embeddings, language models and graphs are components of that chain.

For A2 family-office teams, the engine can support manager diligence, investment-committee memory and portfolio research within strict family and client boundaries. For B2 GCC fund managers, it can support evidence-backed DDQ responses, fundraising consistency and relationship memory within approved disclosure controls. Both audiences require clear decision rights, current claims, traceable calculations, permission isolation and meaningful review.

The implementation rule is concise: **begin with one decision workflow; ingest only authorised sources; retrieve within purpose and permission; cite every material claim; surface conflicts and staleness; abstain when evidence is insufficient; preserve human authority; evaluate severe failures as hard gates; and recognise financial value only after approved observed attribution.**

The reviewed sources support disciplined experimentation and governance. They do not establish Matchpoint or client financial benefit. Attributed revenue, cash cost reduction and loss reduction remain USD 0 until approved observed evidence exists.

REFERENCES

- [1] Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2), 109-122. <https://doi.org/10.1002/smj.4250171110>
- [2] Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14-37. <https://doi.org/10.1287/orsc.5.1.14>
- [3] Alavi, M. and Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems. *MIS Quarterly*, 25(1), 107-136. <https://doi.org/10.2307/3250961>
- [4] Argote, L. and Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes*, 82(1), 150-169. <https://doi.org/10.1006/obhd.2000.2893>
- [5] Cohen, W. M. and Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128-152. <https://doi.org/10.2307/2393553>
- [6] Szulanski, G. (1996). Exploring internal stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management Journal*, 17(S2), 27-43. <https://doi.org/10.1002/smj.4250171105>
- [7] March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71-87. <https://doi.org/10.1287/orsc.2.1.71>
- [8] Kogut, B. and Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3(3), 383-397. <https://doi.org/10.1287/orsc.3.3.383>
- [9] Lewis, P. et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33. <https://papers.nips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>
- [10] Karpukhin, V. et al. (2020). Dense passage retrieval for open-domain question answering. *Proceedings of EMNLP 2020*, 6769-6781. <https://doi.org/10.18653/v1/2020.emnlp-main.550>
- [11] Thakur, N. et al. (2021). BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint arXiv:2104.08663*. <https://arxiv.org/abs/2104.08663>
- [12] Liu, N. F. et al. (2024). Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12, 157-173. https://doi.org/10.1162/tacl_a_00638
- [13] Asai, A. et al. (2024). Self-RAG: Learning to retrieve, generate, and critique through self-reflection. *International Conference on Learning Representations 2024*. <https://openreview.net/forum?id=hSyW5go0v8>
- [14] Es, S. et al. (2024). RAGAS: Automated evaluation of retrieval augmented generation. *Proceedings of EACL 2024: System Demonstrations*, 150-158. <https://doi.org/10.18653/v1/2024.eacl-demo.16>
- [15] Saad-Falcon, J. et al. (2024). ARES: An automated evaluation framework for retrieval-augmented generation systems. *Proceedings of NAACL 2024*, 338-354. <https://doi.org/10.18653/v1/2024.naacl-long.20>
- [16] Edge, D. et al. (2024). From local to global: A graph RAG approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*. <https://arxiv.org/abs/2404.16130>
- [17] Guu, K. et al. (2020). REALM: Retrieval-augmented language model pre-training. *Proceedings of ICML 2020, PMLR 119*, 3929-3938. <https://proceedings.mlr.press/v119/guu20a.html>
- [18] Borgeaud, S. et al. (2022). Improving language models by retrieving from trillions of tokens. *Proceedings of ICML 2022, PMLR 162*, 2206-2240. <https://proceedings.mlr.press/v162/borgeaud22a.html>
- [19] Izacard, G. et al. (2022). Atlas: Few-shot learning with retrieval augmented language models. *arXiv preprint arXiv:2208.03299*. <https://arxiv.org/abs/2208.03299>
- [20] Khattab, O. and Zaharia, M. (2020). ColBERT: Efficient and effective passage search via contextualized late interaction over BERT. *Proceedings of SIGIR 2020*, 39-48. <https://doi.org/10.1145/3397271.3401075>
- [21] Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192. <https://doi.org/10.1126/science.adh2586>

- [22] Brynjolfsson, E., Li, D. and Raymond, L. R. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889-942. <https://doi.org/10.1093/qje/qjae044>
- [23] Dell'Acqua, F. et al. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Organization Science, 37(2), 403-423. <https://doi.org/10.1287/orsc.2025.21838>
- [24] Wiles, E. et al. (2024). GenAI as an exoskeleton: Experimental evidence on knowledge workers using GenAI on new skills. SSRN working paper 4944588. <https://doi.org/10.2139/ssrn.4944588>
- [25] 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>
- [26] 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>
- [27] World Wide Web Consortium (2013). PROV-DM: The PROV Data Model, W3C Recommendation. <https://www.w3.org/TR/prov-dm/>
- [28] National Institute of Standards and Technology (2022). Secure Software Development Framework (SSDF) Version 1.1, NIST SP 800-218. <https://doi.org/10.6028/NIST.SP.800-218>
- [29] OWASP Foundation (2025). OWASP Top 10 for LLM Applications 2025. <https://owasp.org/www-project-top-10-for-large-language-model-applications/>
- [30] MITRE (2026). ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems. <https://atlas.mitre.org/>
- [31] United Arab Emirates Government (2026). Data protection laws. <https://u.ae/en/about-the-uae/digital-uae/data/data-protection-laws>
- [32] Dubai International Financial Centre (2025). Data Protection Law, DIFC Law No. 5 of 2020, as amended. <https://www.difc.com/business/laws-and-regulations/legal-database/difc-laws/data-protection-law-difc-law-no-5-2020>
- [33] Abu Dhabi Global Market (2026). Data Protection Regulations 2021, updated, and official guidance. <https://www.adgm.com/operating-in-adgm/office-of-data-protection/resources>
- [34] Dubai Financial Services Authority (2025). New DFSA AI survey: Generative AI adoption has nearly tripled within the DIFC in the last 12 months; governance continues to develop. <https://www.dfsa.ae/news/new-dfsa-ai-survey-generative-ai-adoption-has-nearly-tripled-within-difc-last-12-months-governance-continues-develop>
- [35] Dubai Financial Services Authority (2026). Regulatory expectations on AI risk management in the DIFC, SEO Letter, 4 June 2026. <https://www.dfsa.ae/your-resources/publications-reports/seo-letters-1>
- [36] U.S. Securities and Exchange Commission (2020). Investment Adviser Marketing: Final Rule, Release No. IA-5653. <https://www.sec.gov/files/rules/final/2020/ia-5653.pdf>
- [37] U.S. Securities and Exchange Commission (2024). Initial observations regarding Advisers Act Marketing Rule compliance, Risk Alert, 17 April 2024. <https://www.sec.gov/compliance/risk-alerts/risk-alert-041724>
- [38] U.S. Securities and Exchange Commission (2024). SEC charges two investment advisers with making false and misleading statements about their use of artificial intelligence, Press Release 2024-36. <https://www.sec.gov/newsroom/press-releases/2024-36>
- [39] International Organization of Securities Commissions (2025). Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges, Consultation Report CR/01/2025. <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf>
- [40] National Institute of Standards and Technology (2020). NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management, Version 1.0. <https://doi.org/10.6028/NIST.CSWP.01162020>

APPENDIX A. SCENARIO ASSUMPTIONS REGISTER

Assumption	Status	Approval required before use
Corpus size and source mix	Unverified illustrative management assumption	Knowledge and system owners
User and reviewer counts	Unverified illustrative management assumption	Business owner
Evaluation-set size	Unverified illustrative management assumption	Validator
Task time and accepted hours released	Unverified illustrative management assumption	Operations and finance
Model, hosting and integration cost	Unverified illustrative management assumption	Technology and finance
Loaded cost and realisation factor	Unverified illustrative management assumption	Finance
Revenue or loss attribution	Unverified illustrative management assumption	Finance and management
Privacy and regulatory treatment	Requires fact-specific advice	Legal, privacy and compliance

APPENDIX B. MINIMUM KNOWLEDGE RECORD

Object	Minimum fields
Source	ID, checksum, owner, dates, rights, client, jurisdiction, retention, parser
Claim	Exact statement, passages, entities, units, status, reviewer, validity
Decision	Question, alternatives, evidence, assumptions, dissent, authority, date
Outcome	Observation, date, evidence, attribution, approved lesson
Relationship	Entities, relationship type, source, date, status and permission
Generated output	Query, purpose, sources, model, workflow, reviewer and approval
Incident	Scope, severity, containment, correction, notification and prevention

APPENDIX C. MINIMUM EVALUATION SET

Test family	Required cases
Retrieval	Exact, semantic, numerical, table, multilingual and supersession
Generation	Single-source, multi-source, conflict, qualification and no-answer
Permissions	Direct, indirect, inferred and tool-mediated cross-client attempts
Security	Injection, poisoning, malicious file, output handling and vendor change

Test family	Required cases
Governance	Approval, expiry, correction, incident and audit reconstruction
Economics	Comparable baseline, quality, capacity, realisation and attribution

APPENDIX D. RED-FLAG REGISTER

Red flag	Required response
No named source owner	Stop ingestion
Client or licence right unclear	Quarantine source
Retrieval applies permission after generation	Stop release
Generated claim lacks passage support	Reject output
Material conflict hidden	Reject output and escalate
External claim has no approver	Prohibit publication
Model or parser changed without regression	Roll back or suspend
Production incident absent from test set	Update validation before next release
ROI uses unobserved hours or unapproved rates	Label unverified and attribute USD 0

APPENDIX E. GLOSSARY

Abstention. A controlled output stating that approved evidence, permission or confidence is insufficient for the requested answer.
Claim. A bounded statement that can be supported, contradicted, reviewed, approved or superseded.
Decision memory. The connected record of a question, alternatives, evidence, assumptions, authority, decision and later outcome.
Dense retrieval. Retrieval using learned vector representations of queries and documents.
Faithfulness. The degree to which an answer's material statements are supported by the supplied evidence.
Hybrid retrieval. A retrieval design combining lexical, dense, metadata, graph or other signals.
Knowledge engine. A governed system connecting sources, claims, entities, decisions, outcomes, permissions and workflows.
Provenance. Information describing the origin, transformation, actors and history of an evidence object or output.
Reranker. A model or rule that reorders retrieved candidates for a query.
Supersession. An explicit relationship showing that a newer approved record replaces an older one for a defined use.