AI for Fund Due Diligence: Building an Evidence Graph from DDQs, Data Rooms and Calls

A Controlled Architecture for Manager Evidence, Source Provenance, Exceptions and Investment-Committee Decisions

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Fund due diligence produces a large and fragmented evidence set. A manager may answer a due-diligence questionnaire, upload policies and reports to a data room, provide performance and portfolio files, respond to follow-up questions and explain exceptions during calls. Reviewers must connect those materials to legal terms, operational controls, investment claims, regulatory records and external evidence. Conventional folder structures and spreadsheets often preserve documents while losing the relationship between a question, the evidence relied upon, the conclusion reached and the decision ultimately taken.

This paper develops a controlled evidence-graph architecture for fund due diligence. The graph treats questions, claims, documents, data points, people, entities, controls, exceptions, decisions and monitoring obligations as linked records. It uses artificial intelligence for bounded tasks such as document classification, extraction, entity resolution, semantic retrieval, cross-document comparison and exception triage. Every material conclusion remains subject to human review, source inspection, conflict assessment and approval. The design combines an evidence graph, a DDQ coverage matrix, a source-confidence ladder, an exception workflow and an investment-committee audit trail.

The framework addresses five recurring problems. First, receipt of a document can be mistaken for proof of a claim. Second, the latest response can overwrite earlier contradictions. Third, AI-generated summaries can detach statements from their source and context. Fourth, diligence findings can disappear between workstreams or reviewers. Fifth, the committee pack can present a conclusion without preserving the evidence state and unresolved issues that informed it. The proposed architecture responds by enforcing provenance, record versioning, materiality, decision rights and monitoring hand-offs.

Any scores, thresholds, examples, timelines and portfolio facts in this paper are hypothetical modelling inputs. They demonstrate a governance method and do not represent an identified manager, fund, investor, offer, solicitation, investment recommendation or legal conclusion. A live implementation requires current legal, regulatory, data-protection, cyber-security, accounting, tax and investment advice in each relevant jurisdiction.

JEL Classification: G11, G23, G24, G28, G32, O33

Keywords: fund due diligence, artificial intelligence, evidence graph, DDQ, data room, manager selection, investment committee, source provenance, private markets, human review

1. Treat diligence as a chain of evidence

Fund due diligence is a decision process conducted through evidence. The investor asks whether a manager has the people, process, controls, incentives, resources and track record to execute the

proposed strategy within agreed risk limits. The answer is rarely contained in one document. It emerges from a chain linking a manager statement to supporting records, independent checks, reviewer judgement, exceptions and conditions.

The ILPA DDQ 2.0 provides a widely used private-markets question structure and a requested-document appendix. Its value extends beyond standardisation. It helps define what should be asked, which supporting materials may be expected and how an investor can compare managers. Yet a completed DDQ is still a manager representation. A strong diligence process tests material statements against constitutional documents, audited accounts, portfolio data, policies, service-provider records, regulatory disclosures, background checks and call evidence.

A folder can show that a policy exists. It does not show which question the policy answers, whether it applies to the proposed vehicle, whether it was in force during the relevant period, whether practice matches the policy or whether another document contradicts it. A spreadsheet can track requests, though it becomes difficult to preserve paragraph-level provenance, versions, entity relationships and reviewer reasoning as the evidence set grows.

An evidence graph solves the relationship problem. Each important object receives a stable identifier and each connection has a type. A claim may be supported by a policy, contradicted by a call transcript and qualified by a side letter. A key person may hold roles across the manager, general partner, investment committee and portfolio companies. A valuation control may be described in the DDQ, governed by a policy and evidenced by committee minutes. The graph retains these distinctions.

Figure 1. Evidence graph linking diligence questions to committee decisions

EACH DECISION SHOULD REMAIN TRACEABLE TO QUESTIONS, SOURCES AND REVIEW

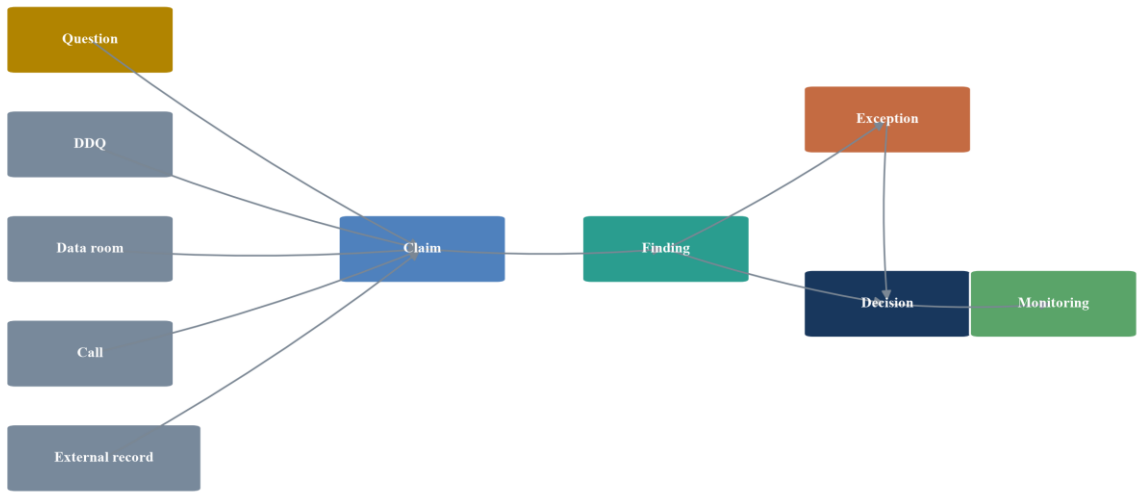

The graph preserves provenance and separates manager representations, source records, reviewer findings and approved decisions.

Table 1. Core evidence-graph objects

Object	Required fields	Typical source	Principal control
question	identifier, topic, materiality, owner, status	approved DDQ or investor request list	controlled question library

Object	Required fields	Typical source	Principal control
claim	exact statement, speaker, entity, period, context	DDQ, document, email or call	preserve wording and attribution
source	file, page, paragraph, date, version, access	data room, register or independent provider	immutable copy and provenance
data point	metric, unit, period, definition, lineage	portfolio file, accounts or system export	validation and reconciliation
finding	reviewer conclusion, evidence set, confidence	analytical workpaper	named reviewer and date
exception	issue, materiality, impact, owner, deadline	contradiction, gap or control failure	escalation and closure evidence
decision	outcome, rationale, conditions, voters	investment-committee record	approval authority and conflicts
obligation	action, trigger, frequency, evidence, owner	side letter, condition or monitoring plan	post-close hand-off and testing

Stable identifiers and dated versions allow the diligence record to be reproduced after approval.

2. Define the diligence ontology before selecting technology

An ontology states what the diligence team needs to know and how concepts relate. Technology selected before the ontology often automates an inconsistent process. The starting point is a question taxonomy covering organisation, ownership, governance, strategy, team, track record, portfolio construction, risk, valuation, operations, compliance, cyber security, service providers, sustainability, legal terms, fees, conflicts and reporting.

The taxonomy should distinguish an entity from a role and a role from a person. A management company, general partner, investment adviser, fund vehicle, carried-interest vehicle and portfolio company may have overlapping owners and directors. A person can be a key person under the fund documents, a voting committee member and an executive of another entity. These relationships matter for governance, conflicts, concentration and succession.

Claims require their own structure. A statement that the manager has never experienced a material cyber incident has a speaker, entity scope, time period, definition of materiality and date. Without those attributes, later evidence cannot be compared fairly. A statement about gross track record should identify strategy, attribution, currency, valuation date, treatment of leverage, fees and realised status. The ontology should store the definition alongside the number.

Controls also need structure. A policy is a design statement. Evidence of operation may include approvals, logs, samples, minutes, testing, incidents and remediation. The graph should therefore connect a control objective to the policy, control owner, operating evidence, test result and exception. This prevents a policy document from being treated as proof that the control operated effectively.

The ontology should remain practical. Every field needs a decision use, reporting purpose or control requirement. Excessive fields create false completeness and slow reviewers. A controlled dictionary, change process and mapping to the DDQ allow the model to evolve without breaking earlier records.

3. Build an ingestion boundary that protects source integrity

The diligence corpus can contain confidential personal data, portfolio-company information, legal advice, investor data and commercially sensitive models. Ingestion therefore begins with authority, classification and access. The team should document which data may be processed, where it may be stored, which model or service may receive it, how long it is retained and how it is deleted.

Each received file should be preserved in its original form with a cryptographic hash, receipt time, source location and access label. A working copy may be converted for search or extraction. The original remains the evidential anchor. Optical character recognition, spreadsheet parsing, transcription and document conversion can introduce errors; derived text should retain a pointer to the page, cell, slide or timestamp from which it came.

Version control is central. Data-room files can be replaced without a visible history. A policy may be uploaded in several versions. An updated DDQ may alter an answer after a challenge. The system should not overwrite earlier evidence. It should record supersession, identify what changed and determine whether prior findings require review.

Calls require informed governance. The team should confirm whether recording or transcription is lawful and agreed, distinguish transcript text from reviewer notes and retain timestamps. Speaker attribution should be verified for material statements. A concise call note can remain the authoritative record when recording is inappropriate.

Third-party AI services create an additional dependency. IOSCO and FCA materials emphasise accountability, skills, oversight and third-party risk. The diligence owner should approve providers, permitted data classes, security controls, contractual terms, logging, model changes, incident response and exit arrangements. Sensitive evidence should not be submitted to a consumer service under an individual account.

Table 2. Controlled ingestion and provenance fields

Field	Purpose	Minimum evidence	Failure prevented
source identifier	stable reference across workstreams	repository record	broken citations and duplicate files
cryptographic hash	prove file identity	calculated at receipt	silent replacement
source party	establish who supplied the item	data-room or correspondence record	uncertain attribution
document date	establish period and currency	document face or metadata	use of stale policy
receipt date	reconstruct evidence available to reviewers	system log	hindsight contamination
version relationship	show replacement or amendment	supersession link	overwritten contradiction
location pointer	return to page, cell or timestamp	extraction metadata	unsupported summary
access class	restrict sensitive content	approved classification	excessive disclosure
processing authority	document permitted AI use	legal and security approval	uncontrolled processing

Field	Purpose	Minimum evidence	Failure prevented
retention rule	manage deletion and legal hold	approved schedule	indefinite accumulation

Derived content remains linked to the immutable source and its access conditions.

4. Map DDQ questions to evidence requirements

A DDQ coverage matrix should answer more than whether a response was received. For each question, it should identify the expected evidence, manager answer, source coverage, reviewer conclusion, exceptions and follow-up. A complete response can remain unsupported; a document can be relevant without addressing the question; and an answer can be internally consistent while conflicting with independent records.

The matrix begins with materiality. Questions affecting legal existence, regulatory status, ownership, key-person dependence, track-record integrity, valuation, custody, cash control, conflicts, sanctions, cyber resilience and fund terms usually require direct review. Lower-risk administrative items can use sampled verification. Materiality should be approved and capable of adjustment when a new fact changes the risk.

Coverage should use separate states. Received means the requested item arrived. Responsive means it addresses the question. Supported means the source provides evidence for the statement. Current means the evidence covers the relevant period. Consistent means it agrees with other material evidence. Reviewed means an accountable professional inspected the source and recorded a conclusion. These states should never collapse into one percentage.

The matrix can also identify evidence dependency. Several answers may rely on the same policy or service-provider report. If that source is stale or out of scope, multiple questions become affected. The graph makes this dependency visible and allows the team to update all connected findings when new evidence arrives.

Figure 2. Illustrative DDQ coverage matrix

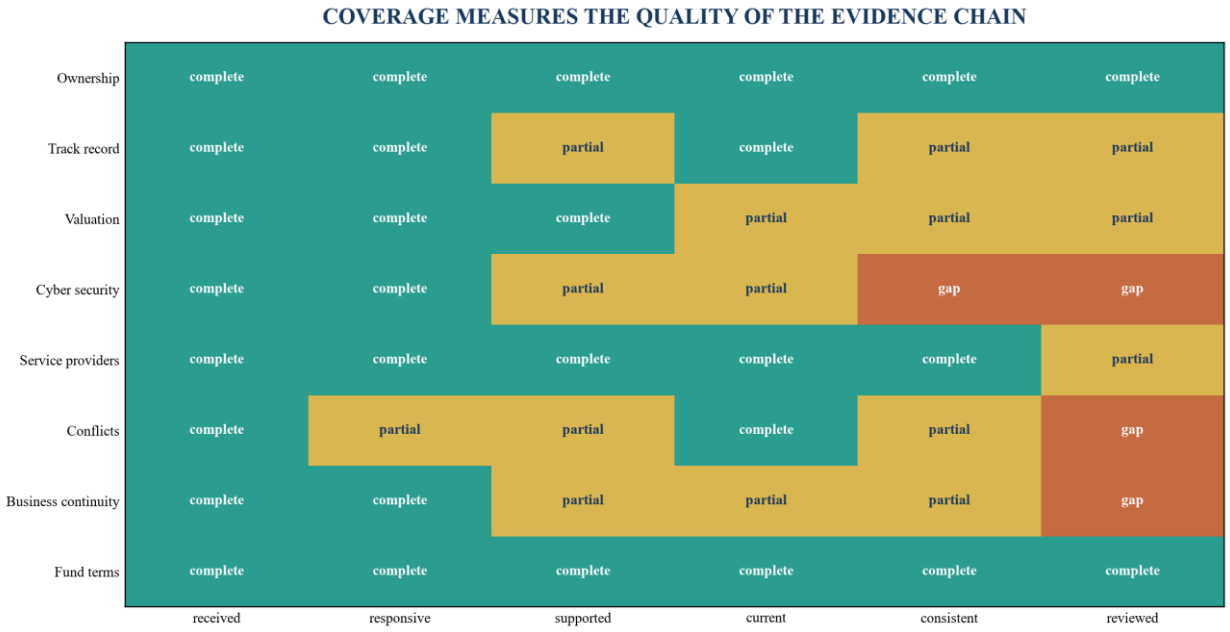

A green receipt indicator does not establish support, currency, consistency or completed human review.

Table 3. Coverage states and required reviewer action

State	Test	Evidence	Reviewer action
received	requested item is present	repository receipt	confirm file identity and scope
responsive	item addresses the question	cited passage or data field	reject generic or unrelated material
supported	statement is evidenced	direct record or independent source	classify strength and limitations
current	period remains relevant	date, effective period and events	request refresh or explain reliance
consistent	evidence agrees across sources	comparison and exception results	investigate contradictions
reviewed	accountable person completed challenge	signed workpaper	record conclusion and open issues
approved	authority accepted residual risk	committee or delegated approval	preserve conditions and rationale
monitored	post-close obligation is active	monitoring plan and evidence cycle	test, escalate and report

Each state is independently testable and can be reported by materiality and workstream.

5. Use AI for bounded evidence tasks

AI is useful when the task is repetitive, evidence is voluminous and the output can be checked. Classification can route a document to a workstream. Extraction can identify names, dates, fees, commitments, service providers, controls and defined terms. Semantic retrieval can surface passages relevant to a question. Comparison can highlight changed DDQ answers or policy clauses. Entity resolution can propose links among people and vehicles. Exception triage can rank missing or contradictory evidence for review.

The system should separate retrieval from conclusion. A model can locate passages discussing valuation conflicts. It should not decide that conflicts are adequately controlled without an approved test and human judgement. A model can calculate whether track-record rows reconcile to totals. It should not determine attribution quality when deal roles, predecessor rights or strategy definitions remain disputed.

Prompts and schemas should be versioned. Each extraction should state the model, configuration, prompt version, time, source identifiers and confidence or validation result. Material outputs need citations at page, paragraph, cell or timestamp level. Unsupported model prose should be excluded from the evidence graph.

Deterministic checks should be used where possible. Dates, totals, duplicate identifiers, missing columns, formula errors and cross-footing are often better tested with code than generated text. AI can explain anomalies after the calculation is preserved. This division improves repeatability and reduces the risk that fluent language hides a numerical error.

Human reviewers should receive the source passage, surrounding context, proposed extraction, reason for the match and any conflicting evidence. Review actions should be quick to record: accept, amend, reject, escalate or request evidence. The system learns operationally through controlled rules and templates, not through invisible changes to an ungoverned model.

6. Grade source confidence without creating false precision

Source confidence is a structured judgement about how much weight a source deserves for a particular claim. It is not a general reputation score. An audited financial statement can provide strong evidence of historical balances within its scope. It may provide little evidence about current team capacity. A call with the chief operating officer can explain a control, while operating samples and independent assurance provide stronger evidence that it worked.

The confidence ladder should assess authority, directness, currency, completeness, consistency and susceptibility to conflict. Independent regulatory or legal records may sit high for status and ownership. Executed fund documents are authoritative for contractual terms. Audited accounts are strong for audited assertions. System exports and samples can support operating evidence when lineage is verified. Policies describe design. DDQ responses and calls are representations that require corroboration according to materiality.

A numeric score can aid triage, but the underlying reasons must remain visible. Two sources with the same score can have different weaknesses. The model should store a confidence class, rationale and reviewer. The class can change when a source expires, an exception arises or independent evidence arrives.

Figure 3. Source-confidence ladder for fund due diligence

Confidence is claim-specific and combines source authority with directness, currency, completeness and consistency.

7. Preserve contradictions and changes as evidence

Contradictions are valuable diligence signals. The team size in a presentation may differ from the regulatory filing. A DDQ can state that valuation is independent while committee minutes show investment professionals approving marks. A policy may require annual testing while the latest report is older. A call can qualify a broad answer given earlier.

The system should preserve each statement, source and effective date. It should create an exception link rather than selecting the latest text and deleting the conflict. The exception record should describe the issue neutrally, show the sources, state why the difference matters, identify the owner and record the resolution standard.

Some differences are legitimate. Entity scope, definition, period or materiality can explain them. Resolution should therefore identify whether the claims were reconciled, one source superseded another, the difference was immaterial, a remediation was accepted or the issue remains open. The decision record should show which conclusion the investor used and why.

AI can detect semantic differences that exact matching misses, including changes in qualifiers such as normally, subject to approval or where practicable. It can also create false conflicts when two statements use different scopes. Human review is essential before escalation. The original passages and context should be displayed together.

8. Create an exception workflow with defined decision rights

An exception is a material gap, contradiction, control weakness, unexplained result or evidence limitation. Every exception needs a type, severity, affected claims, potential decision impact, owner, response date and closure test. A missing document can be low severity if another reliable source addresses the matter. A small numerical difference can be high severity when it indicates an unreliable track-record process.

The workflow should distinguish evidence requests, analytical review, risk acceptance, remediation and committee conditions. Reviewers can close an administrative gap when the required item arrives. A material conflict may require legal, compliance, operational, tax, valuation or investment specialists. Acceptance of residual risk belongs to the authorised decision body.

Deadlines should reflect decision dependency. An issue that prevents assessment should be resolved before recommendation. A remediable weakness may become a condition precedent, side-letter obligation or monitoring action. The committee should see what remains open, who can accept it and what happens if the condition is missed.

Figure 4. Exception workflow from detection to monitoring

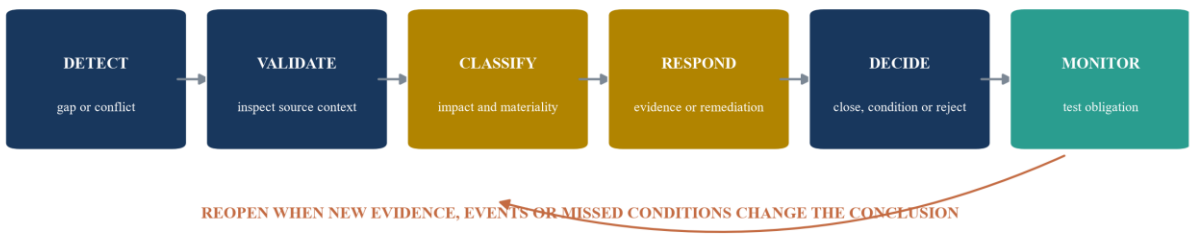

Closure requires evidence and authority; a manager response alone does not close a material exception.

Table 4. Illustrative exception classes and escalation

Exception class	Example	Initial owner	Closure evidence	Decision authority
evidence gap	missing current insurance or policy	operations reviewer	verified current document	workstream lead
inconsistency	team count differs across records	investment reviewer	scoped reconciliation and source	diligence lead
control design	policy lacks independent approval	operational diligence	approved revised design	risk or committee delegate
operating failure	required review was missed	control specialist	root cause, remediation and test	investment committee
legal or regulatory	entity status or term is uncertain	counsel or compliance	authoritative record and advice	authorised committee
financial	track record or fee calculation does not reconcile	finance reviewer	corrected lineage and validation	investment committee
conduct or integrity	disclosure appears incomplete or misleading	senior diligence owner	independent investigation	governing body
post-close condition	remediation cannot finish before approval	obligation owner	dated deliverable and test	committee under documented terms

Materiality depends on the strategy, investor, jurisdiction and decision context.

9. Test track record and portfolio data through lineage

Track-record analysis needs a path from reported fund and deal metrics to underlying records. Each observation should identify the investment, vehicle, ownership period, cash-flow date, currency, gross or net basis, valuation date, attribution rights and source. Derived returns should preserve the calculation and inputs.

AI can assist with mapping inconsistent investment names, classifying cash flows, comparing presentation tables with detailed files and locating narrative explanations. Deterministic calculations should reproduce totals and returns. Reviewers should test definition changes, predecessor performance, partial realisations, subscription facilities, foreign exchange, write-offs, continuation vehicles and unrealised value.

Portfolio data should also connect to claims about strategy. If a manager describes lower-middle-market control investing, the graph can test entry enterprise value, ownership, geography, sector and transaction type. If the manager claims operational value creation, the investor can link initiatives to portfolio evidence and separate market, leverage, multiple and earnings effects where data permit.

Sampling should be risk-based. Large winners, write-offs, valuation outliers, related-party transactions, amended deals and investments near reporting cut-offs deserve attention. The evidence graph should record why each sample was selected and whether an issue affects the wider population.

10. Connect operational due diligence to control evidence

Operational due diligence examines whether the manager can protect assets, calculate values, control cash, maintain records, comply with obligations and continue operating through disruption. The evidence graph should link each control objective to its owner, policy, system, service provider, operating evidence, assurance, incidents and remediation.

Cash control illustrates the method. The manager may describe dual authorisation in the DDQ and treasury policy. The reviewer should identify bank accounts, signatories, payment initiation, approval thresholds, callback procedures, access reviews, segregation, reconciliations, exceptions and incident history. Samples and logs test operation. Service-provider reports may add evidence, subject to scope, period and user-control considerations.

Outsourcing does not transfer accountability. FCA guidance states that firms remain responsible for managing risks arising from outsourcing and third-party arrangements. IOSCO's AI guidance likewise emphasises designated responsibility, adequate skills, testing, monitoring and provider oversight. The graph should show the dependency chain from the manager to administrators, custodians, technology providers, data vendors and sub-processors.

Business continuity evidence should cover important services, impact tolerances, people, facilities, technology, information and third parties. Plans should be connected to tests, findings and remediation. A successful annual exercise has limited value if it omitted a critical provider or did not test a severe scenario. The reviewer should record what was tested and what remains outside scope.

Cyber evidence needs controlled access. Architecture diagrams, penetration-test findings and incident records can be highly sensitive. The investor can record a conclusion, evidence identifier, scope, date and reviewer without exposing raw material broadly. Access controls should follow the sensitivity of the source, while the committee receives the decision-relevant finding and limitations.

11. Evaluate ownership, governance and conflicts as relationships

Ownership and governance are naturally graph-shaped. The manager, holding company, general partner, funds, carried-interest vehicles, affiliates and portfolio companies may share owners, directors, employees and service providers. A relationship model can reveal conflicts and dependencies that a flat questionnaire obscures.

Beneficial ownership should be supported by authoritative corporate records, constitutional documents, cap tables and confirmations appropriate to the jurisdiction. FATF's strengthened standards emphasise adequate, accurate and current beneficial-ownership information for legal persons and arrangements. The diligence team should identify the natural persons who ultimately own or control relevant entities and document any limitation in the available evidence.

Governance analysis should map formal authority and practical influence. Committee terms, voting rules, quorum, reserved matters, conflicts procedures and minutes establish design and operation. The graph can compare the stated investment process with the decisions recorded in samples. It can identify when one person appears across origination, valuation, risk and final approval, creating concentration or independence concerns.

Conflicts should be attached to transactions and decisions. Allocation among funds, cross trades, continuation vehicles, affiliated services, financing arrangements, co-investments, personal investments, valuation and fees can create different incentives. A conflict register alone does not establish effective management. The reviewer should connect the conflict to disclosure, approval, pricing, recusal, investor consent and operating evidence.

AI can support entity resolution across records, subject to review. Similar names, transliteration, holding companies and common addresses can produce false matches. Each proposed link should display the matching attributes and confidence. Material ownership conclusions require authoritative evidence and human approval.

12. Integrate legal terms with commercial findings

Legal due diligence and commercial diligence should share a controlled issue map. A commercial concern about team concentration may be addressed partly through key-person provisions, though the legal definition, trigger, remedy and replacement process determine protection. A valuation concern may connect to advisory-committee rights, reporting, audit and removal provisions. A conflict concern may connect to consent, disclosure and allocation clauses.

The evidence graph should store each material term with its source clause, defined terms, affected entity, condition and reviewer interpretation. It should link the term to the underlying risk and any side-letter request. This enables the committee to see whether a risk is avoided, controlled, priced, disclosed, accepted or monitored.

AI can retrieve clauses and compare documents, but legal meaning depends on the entire instrument, governing law and facts. Extracted terms should retain the surrounding clause and be validated by qualified counsel. A model-generated summary should never become the authoritative legal record.

Negotiation changes need version control. The system should compare drafts, show which risk a change addresses and confirm the executed document contains the approved term. Conditions negotiated in correspondence can be lost if they are not linked to closing deliverables and the monitoring plan.

13. Build the investment-committee audit trail

The investment committee needs a concise decision pack and a reproducible evidence record. The pack should state the proposed commitment, strategy fit, expected return and risk, key strengths, material exceptions, legal protections, conditions, monitoring plan and recommendation. Each material statement should link to an approved finding rather than directly to raw model output.

The audit trail should preserve the evidence state at the decision cut-off. Data-room content can change, a manager can provide later clarification and market conditions can move. The record should identify which sources, versions and findings were available, which issues remained open and which assumptions the committee accepted.

Decision authority and conflicts must be explicit. The record should identify attendees, voting members, recusals, quorum, votes, delegated matters and conditions. A committee decision may

approve, reject, defer, request further evidence, reduce the commitment or impose conditions. The rationale should explain the link between evidence and outcome.

Post-meeting actions should become obligations with owners, dates, evidence and escalation. Approval subject to a side letter, policy revision, key hire or reporting deliverable is incomplete until the condition is satisfied. The graph should prevent closure when evidence is missing and reopen the issue when a later event changes the basis of the decision.

Figure 5. Investment-committee audit trail

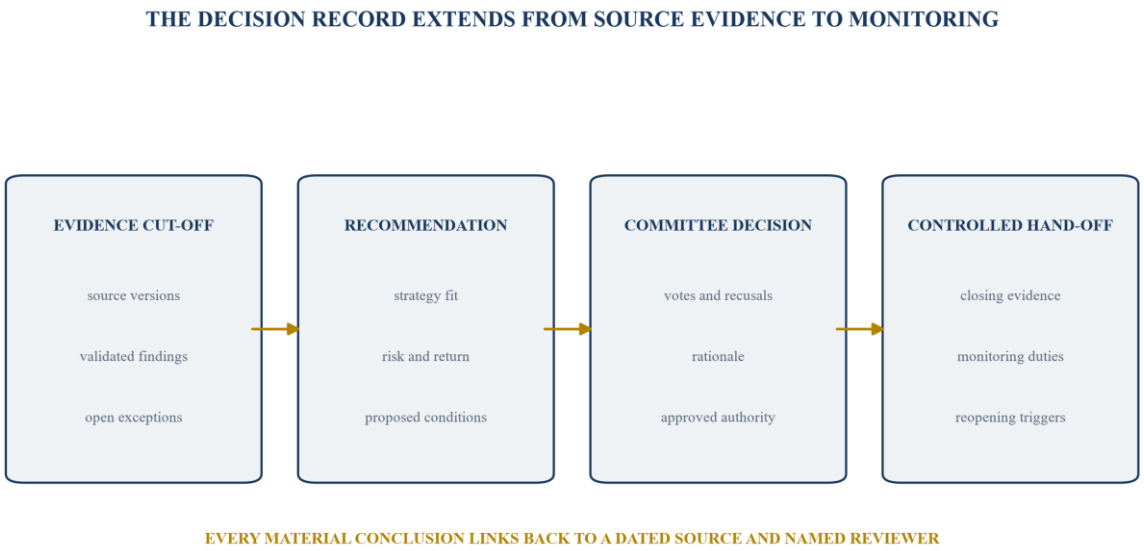

The record captures the evidence cut-off, decision rationale, conditions and post-close obligations.

Table 5. Investment-committee record and control owner

Record component	Required content	Owner	Verification
evidence cut-off	date, repository snapshot and source versions	diligence lead	manifest and hashes
findings	approved conclusions, evidence and confidence	workstream reviewers	reviewer sign-off
exceptions	open, closed, accepted and conditional issues	exception owners	status and closure evidence
recommendation	commitment, rationale, alternatives and risks	sponsor or investment team	pack approval
conflicts	member, interest, recusal and impact	committee chair or secretary	attendance and voting record
decision	approve, reject, defer, limit or condition	authorised committee	minutes and authority
conditions	action, deadline, evidence and consequence	named obligation owner	close checklist
monitoring	indicator, frequency, trigger and escalation	portfolio or monitoring team	scheduled evidence test

The record should be retained in line with applicable law, policy and contractual obligations.

14. Govern models, prompts and human review

The AI components of the diligence system require their own governance. NIST's AI Risk Management Framework organises activity around govern, map, measure and manage. Its Generative AI Profile highlights risks including confabulation, data privacy, information integrity, security, intellectual property, harmful bias, human over-reliance and third-party dependencies. These risks translate directly into diligence controls.

Governance should begin with an inventory of models and use cases. Each use case needs an owner, permitted data, decision impact, validation method, access model, provider, change process and retirement plan. A classifier that routes documents has a different risk from a system that drafts investment conclusions. The latter should face stronger restrictions and review.

Testing should use representative documents, difficult layouts, tables, scanned files, multiple languages, ambiguous entities and known contradictions. Metrics can include extraction accuracy, citation accuracy, missed material issues, false exceptions and reviewer override rates. Results should be segmented by document type and use case. A high average can conceal failure on the records that matter most.

Human review must be substantive. An interface that asks a reviewer to approve hundreds of fluent summaries can create automation bias. Reviewers need source context, clear reasons for escalation and enough time to challenge. The process should monitor rubber-stamping, disagreement, error patterns and workload.

Model and prompt changes should trigger proportionate revalidation. A provider can change a hosted model or safety layer. The firm should know when a material component changed and whether earlier performance assumptions remain valid. Logs should preserve the model and configuration used for material outputs.

15. Secure the evidence system and its data

The evidence graph concentrates valuable information. It can contain identities, ownership, investment records, legal documents, cyber findings and committee decisions. Security architecture should apply least privilege, strong authentication, encryption, logging, environment separation, data-loss controls, backup and tested recovery.

Access should follow workstream and need. A tax adviser may need specific structures and records. A cyber specialist may need sensitive technical evidence. Committee members need decision-relevant findings and access to supporting sources under controlled permissions. Graph links should respect the access level of every connected object and avoid leaking restricted information through search snippets or generated summaries.

Prompt injection and malicious documents require attention. A document can contain text designed to influence an AI system. The pipeline should treat document content as evidence, not instruction; isolate tools; restrict actions; validate outputs; and require approval before any external communication or system change. Retrieval should never grant a model authority to send requests, amend records or approve findings.

Incident response should cover incorrect extraction, unauthorised access, provider breach, data leakage, model failure and corrupted provenance. The plan should identify containment, evidence preservation, legal and regulatory assessment, notification, remediation and revalidation. The

committee should understand any incident that materially affects the integrity of its decision record.

16. Design operating roles and service levels

An evidence graph changes work allocation. The diligence sponsor owns the investment question and recommendation. Workstream specialists own their findings. A data or knowledge team maintains the ontology, ingestion and lineage. Model-risk or technology-risk personnel validate AI components. Legal, compliance, privacy and cyber teams approve relevant boundaries. The investment committee accepts residual investment risk within its authority.

A RACI matrix should identify who requests evidence, validates extraction, resolves exceptions, approves scope changes and signs the final pack. The system should prevent a technology administrator from changing a reviewer conclusion and prevent a reviewer from altering source records. Emergency access should be logged and reviewed.

Service levels should prioritise materiality. High-risk exceptions and missing gating evidence need rapid escalation. Lower-risk classification errors can enter a batch correction. The dashboard should show aging, reviewer capacity, evidence dependencies and decision deadlines without encouraging teams to close issues merely to improve a completion score.

Training should cover the diligence methodology, source hierarchy, system use, confidentiality, model limitations, reviewer accountability and escalation. Reviewers need enough technical understanding to challenge outputs and enough investment context to recognise when a source is decision-relevant.

17. Implement through controlled stages

Implementation should begin with a narrow use case and a completed diligence file. The team can map the question taxonomy, ingest sources, reproduce known findings and measure whether the graph improves retrieval, coverage and auditability. Historical testing reveals gaps without placing a live decision at risk.

The first stage establishes source preservation, identifiers, access and citations. The second adds question and claim mapping. The third introduces deterministic validation and bounded AI extraction. The fourth connects exceptions and workstream approvals. The fifth generates committee views from approved findings. Monitoring obligations follow once the pre-investment record is stable.

Each stage needs acceptance criteria. Source pointers should return to the correct page or cell. Material extractions should meet approved accuracy thresholds by document type. Contradictions should not be silently overwritten. Restricted evidence should remain protected through search and export. Committee output should exclude unreviewed findings.

Parallel operation may be necessary. The existing diligence tracker remains authoritative until the new system demonstrates completeness, reliability and user adoption. Reconciliation between systems should be documented. Migration should preserve historical dates, authors, decisions and evidence rather than converting everything into a current-state summary.

The final gate is operational resilience. The team should test provider failure, loss of a model, corrupted extraction, inaccessible data-room links, staff absence and urgent committee timing. A

diligence process should remain capable of reaching a controlled decision when an AI component is unavailable.

18. Measure whether the system improves decisions

Success should be measured through evidence quality, decision control and operating efficiency. Useful indicators include the percentage of material findings with paragraph-level provenance; material questions supported by current evidence; exceptions resolved with closure evidence; committee conditions handed into monitoring; reviewer override rates; citation errors; and time spent locating sources.

Speed alone is a weak measure. A faster pack can be worse if unsupported summaries pass through review. The system should measure missed material issues, false confidence, unreviewed exceptions and rework. Quality sampling by an independent reviewer can compare conclusions with sources and assess whether the stated confidence is justified.

Decision outcomes require caution. Fund performance is influenced by markets, strategy and events beyond diligence. The organisation can still test whether identified risks were monitored, whether conditions were met, whether later incidents were foreshadowed in evidence and whether the original rationale remains reproducible.

An effective evidence graph creates institutional memory. It allows the investor to compare representations across funds and vintages, identify recurring control weaknesses, update findings when a service provider or person changes and carry pre-investment conditions into ownership. The commercial value comes from faster access to reliable evidence, stronger challenge and a more defensible decision process.

The governing principle is simple: AI may organise, retrieve, compare and propose; accountable professionals inspect sources, resolve uncertainty and decide. The evidence graph gives those professionals a structured record that remains usable after the data room closes and the committee meeting ends.

References

[1] Institutional Limited Partners Association, Due Diligence Questionnaire 2.0, 1 November 2021. https://ilpa.org/resources-tools/resource-library/due-diligence-questionnaire/

[2] Institutional Limited Partners Association, ILPA DDQ 2.0, November 2021. https://ilpa.org/wp-content/uploads/2021/11/ILPA-DDQ-2.0.pdf

[3] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, January 2023. https://doi.org/10.6028/NIST.AI.100-1

[4] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, July 2024, updated April 2026. https://doi.org/10.6028/NIST.AI.600-1

[5] National Institute of Standards and Technology, AI Risk Management Framework Playbook, current resource. https://airc.nist.gov/airmf-resources/playbook/

[6] International Organization of Securities Commissions, The Use of Artificial Intelligence and Machine Learning by Market Intermediaries and Asset Managers, September 2021. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf

[7] International Organization of Securities Commissions, Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges, March 2025. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf

[8] International Organization of Securities Commissions, Supervisory Toolkit for AI Use in Capital Markets, May 2026. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD810.pdf

[9] European Securities and Markets Authority, Public Statement on the Use of Artificial Intelligence in the Provision of Retail Investment Services, 30 May 2024. https://www.esma.europa.eu/document/public-statement-ai-and-investment-services

[10] U.S. Securities and Exchange Commission, Observations from Examinations of Private Fund Advisers, 27 January 2022. https://www.sec.gov/compliance/risk-alerts/observations-examinations-private-fund-advisers

[11] U.S. Securities and Exchange Commission, Risk Alert: Investment Adviser Due Diligence Processes for Selecting Alternative Investments, 28 January 2014. https://www.sec.gov/newsroom/press-releases/2014-14

[12] U.S. Securities and Exchange Commission, Form ADV and Investment Adviser Registration resources, current. https://www.sec.gov/investment/registration-forms-investment-advisers

[13] Financial Conduct Authority, Outsourcing and Operational Resilience, updated 14 July 2026. https://www.fca.org.uk/firms/outsourcing-and-operational-resilience

[14] Financial Conduct Authority, SYSC 8.1 General Outsourcing Requirements, current FCA Handbook. https://handbook.fca.org.uk/handbook/sysc8/sysc8s1

[15] Financial Conduct Authority, Regulatory Priorities: Wholesale Buy Side, March 2026. https://www.fca.org.uk/publication/regulatory-priorities/wholesale-buy-side-report.pdf

[16] Financial Action Task Force, Beneficial Ownership, current standards and guidance. https://www.fatf-gafi.org/en/topics/beneficial-ownership.html

[17] Financial Action Task Force, The FATF Recommendations, updated February 2025. https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Fatf-recommendations.html

[18] Financial Action Task Force, Guidance on Beneficial Ownership of Legal Persons, March 2023. https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Guidance-Beneficial-Ownership-Legal-Persons.html

[19] International Organization for Standardization, ISO/IEC 42001:2023 Information technology, Artificial intelligence, Management system, current overview. https://www.iso.org/standard/81230.html

[20] Bank for International Settlements, The Use of Artificial Intelligence for Policy Purposes, 10 October 2025. https://www.bis.org/publ/othp100.htm

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds, bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.