

AI for LP Operations: Automating Requests while Preserving Accountability

A Controlled Operating Model for Investor-Service Workflows

Authored by Chennakeshav Adya

Independent Researcher

August 2026

Abstract

Limited partner operations contain a large volume of repetitive, document-intensive work: due diligence questionnaires, recurring data requests, investor reports, capital notices, side-letter questions, data-room updates and evidence searches. Artificial intelligence can help classify requests, retrieve approved information, compare templates, identify missing evidence and prepare draft responses. The same workflow can also expose confidential data, lose source lineage, reproduce stale answers, misread fund boundaries, overlook contractual differences or create an investor communication without accountable review.

This paper develops a controlled operating model for applying artificial intelligence to limited partner operations. It separates intake, retrieval, calculation, drafting, review, approval and release; classifies use cases by consequence; grounds generated content in approved sources; preserves citations and records; and assigns accountable human decision rights. The framework draws on the National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework and Generative AI Profile, current Financial Conduct Authority accountability materials, the European Union Artificial Intelligence Act, United Kingdom Information Commissioner's Office data-protection guidance, United States Securities and Exchange Commission books-and-records materials, Financial Stability Board work on artificial intelligence vulnerabilities, and Institutional Limited Partners Association reporting standards.

Six original figures and six implementation tables support an approved-answer library, evidence graph, human-review matrix, testing programme, ninety-day implementation and recurring AI operations office. Every worked request volume, accuracy rate, response time, staffing assumption and timetable is a management assumption used solely to demonstrate the method. These assumptions are not actual operating results, forecasts, regulatory conclusions or evidence of investor demand. Live use requires current legal, regulatory, data-protection, cybersecurity, records-management, contractual and fund-specific review in every relevant jurisdiction.

JEL Classification: G23, G28, O33, L86

Keywords: artificial intelligence, limited partner operations, investor relations, private funds, due diligence, reporting, accountability

1. Define the operating problem before selecting a model

LP operations begin with obligations and decisions, not a technology. A manager receives questions from existing investors, prospective investors, consultants, administrators, auditors, regulators and internal teams. The work can require information about strategy, performance, fees, portfolio companies, governance, sustainability, incidents, side letters, capital activity and legal terms. Each answer can carry a different consequence.

The first control is a request register. Every item receives an investor, fund, entity, topic, channel, due date, confidentiality level, legal or contractual basis, source owner, response owner, approver and release status. Repeated questions become visible. So do conflicts between investor templates, reporting periods and definitions.

Artificial intelligence can support a bounded step within this process. It can classify an incoming request, locate related approved answers, compare a new questionnaire with prior versions, identify source gaps, draft a response from cited evidence or check a completed package for consistency. The use case should be described in operational terms, including its permitted input, output and decision boundary.

A model should not become the system of record. The request register, evidence repository, calculation files, answer library, approvals and released communications remain controlled records. This makes it possible to replace a model or vendor without losing operating memory.

Figure 1. Controlled LP-request workflow
AUTOMATION SITS INSIDE THE OPERATING CONTROL CHAIN



Every released answer retains its fund boundary, sources, reviewer and approval

Artificial intelligence supports bounded tasks within an accountable operating chain.

2. Map the request universe

The operating map should cover due diligence questionnaires, requests for proposals, operational diligence, quarterly and annual reporting, portfolio metrics, tax and regulatory information, capital calls, distributions, statements, audit confirmations, most-favoured-nation elections, side-letter obligations, advisory committee materials, incident notices and ad hoc investor questions.

These workflows differ in structure. A standard template can have defined fields and calculation rules. An ad hoc question can require interpretation. A capital notice can trigger movement of cash. A side-letter response can affect contractual rights. A performance answer can create marketing and recordkeeping consequences. Automation controls therefore follow the use case rather than applying one policy to all investor communication.

The request map records source systems and handoffs. Fund administrator, finance, investment, compliance, legal, tax, sustainability and portfolio teams can own different fields. A missing ownership map causes the AI layer to retrieve the nearest available document rather than the

authoritative one.

The manager should also identify duplicate work. Several LPs may request the same concept through different wording. Standardisation initiatives such as ILPA's DDQ, Reporting Template, Performance Template and updated 2025 Capital Call and Distribution Template can support consistent definitions and transaction reporting. Investor-specific terms and current fund documents still govern the live response.

Table 1. LP operations use-case register

Use case	Typical input	AI-supported task	Accountable output
DDQ	questionnaire and fund evidence	classify, retrieve and draft	approved investor response
recurring report	ledger, valuations and narrative	populate, compare and check	signed reporting pack
capital notice	administrator and fund records	field validation and variance check	authorised notice
side-letter request	agreement and obligation register	locate and compare clauses	legal or compliance response
data-room update	documents and index	classify, tag and identify gaps	approved room release
ad hoc question	investor email and context	route and prepare cited draft	reviewed communication
incident communication	issue record and obligations	identify recipients and draft	authorised notification

Consequence and approval depend on facts, agreements and applicable requirements.

3. Classify use cases by consequence


Risk classification begins with the harm that a wrong, incomplete, late or exposed output could cause. The taxonomy can consider financial movement, investor rights, regulatory communication, performance presentation, personal data, confidential portfolio information, public disclosure and operational reversibility.

Low-consequence assistance can include document tagging, meeting transcription within an approved environment, search across public policies and formatting a previously approved answer. Medium-consequence work can include first drafts of DDQ responses and reconciliation of reporting templates. High-consequence work can include capital notices, contractual interpretations, performance calculations, regulated disclosures, incident notifications and messages that can change investor rights.

The classification sets permitted automation. Low-consequence work may allow automated processing with sampled review. Medium-consequence work requires source citations and named review. High-consequence work requires deterministic calculations or authoritative systems, specialist review, dual approval and controlled release. Certain uses can remain prohibited because confidentiality, legality or reliability cannot be supported.

The EU AI Act applies according to roles, system characteristics and scope. It includes AI literacy obligations and specific requirements for defined high-risk systems and certain transparency situations. A private-fund workflow should not be assigned an EU AI Act category by analogy alone. Current legal analysis is needed where the Regulation applies.

Figure 2. Consequence-based automation ladder
HIGHER CONSEQUENCE REQUIRES STRONGER EVIDENCE AND HUMAN AUTHORITY



The illustrated controls are a management framework; applicable legal classifications require separate review.

4. Establish the source hierarchy

AI output quality begins with source authority. The hierarchy should identify executed fund documents, signed side letters, administrator records, audited statements, approved valuation materials, board or committee records, compliance policies, released investor reports, approved answer components and working drafts.

Documents need fund, entity, investor, period, status, version, owner, confidentiality, effective date and supersession metadata. A search system that retrieves a valid answer for the wrong fund or period can generate a persuasive error. Retrieval therefore filters on authority and perimeter before semantic similarity.

Ingestion is a controlled process. The system verifies file type, malware status, document identity, completeness, effective date and access rights before indexing. Optical character recognition output is compared with the original where a poor scan can change a figure, clause or table. Embedded instructions inside a retrieved document are treated as content rather than commands to the AI system.

Tables, footnotes, schedules and amendments require particular care. A clause may depend on a defined term in another section or an amendment executed later. A data table can lose headers when split into retrieval chunks. The ingestion record therefore preserves document structure and links related provisions. Contractual retrieval can use exact references and specialist review rather than isolated semantic passages.

The system should preserve exact evidence references. A drafted response cites the file, section, page or structured field used. Calculated values cite the controlled calculation output rather than a narrative document. When evidence conflicts, the system creates an exception and stops the affected answer.

Public regulatory and industry sources require version and access dates. A proposal, consultation, final rule and withdrawn proposal have different status. The SEC withdrew its proposed predictive-data-analytics rulemaking in June 2025. A control library should therefore avoid describing that proposal as a current final obligation.

Table 2. Source authority and retrieval controls

Source tier	Examples	Permitted use	Retrieval control
authoritative record	executed agreement, administrator ledger	direct evidence within scope	exact entity, fund, period and version
approved output	signed report or released answer	reusable component	expiry and audience check
controlled calculation	versioned workbook or system result	numerical response	methodology and approval link
policy and governance	approved policy, minutes or register	practice and oversight response	effective-date check
working material	draft memo or email	locate potential evidence	cannot support release alone
external authority	regulator or standard setter	rule and practice context	status and jurisdiction check
model output	generated text or extraction	draft or exception signal	never treated as source evidence

The live hierarchy should be approved by fund, legal, compliance and records owners.

5. Build an approved-answer library

The answer library stores reusable evidence components, not a single universal response. Each component includes topic, direct answer, fund perimeter, period, sources, calculations, limitations, owner, reviewer, approval date, expiry and permitted channels.

Questions are mapped to concepts. Several formulations of a key-person question can retrieve the same approved governance component. The drafting system can assemble fund-specific facts around that component, preserving the original language where the answer is contractual or prescribed.

The library distinguishes stable and volatile information. Firm history, ownership and policy can change infrequently. assets under management, portfolio counts, performance, exposures and personnel can change frequently. Every field receives a refresh rule and event trigger.

An answer that requires a calculation should call a controlled data service or approved output. The language model should not recompute performance, fees, capital accounts or allocations from unstructured documents. It can explain a verified result after retrieving the corresponding method and disclosure.

6. Ground every draft in retrieved evidence

Retrieval-augmented generation can narrow the model's information to approved documents. Grounding requires more than adding a document repository. The retrieval path needs identity filters, access permissions, authority ranking, chunk boundaries, citation capture and an abstention rule.

The system first parses the request into investor, fund, period, question, response type and required fields. It retrieves structured data and documents within that perimeter. It ranks evidence by authority and recency. It then drafts only from cited items. Unsupported required fields remain blank or become exceptions.

The response packet contains draft text, source citations, retrieved passages, calculation identifiers, confidence or evidence-strength flags and detected conflicts. The reviewer can inspect the evidence without repeating the entire search.


The prompt template should state the permitted task, required output schema, source-only rule, fund and period, prohibited content, abstention behaviour and escalation conditions. These controls are versioned and tested. They do not replace model and retrieval evaluation.

The system should make uncertainty operational. A missing required field creates an exception with an owner. Conflicting sources create a hold. A low-authority source can support a search lead while remaining insufficient for release. These states are more useful than a single confidence percentage that reviewers can misinterpret.

Generated citations should be validated. A correct file name does not establish that the cited passage supports the claim. The test checks the exact proposition, entity, time period and amount. Quotations remain short and accurate where used.

Figure 3. Evidence-grounded response architecture

THE MODEL DRAFTS FROM A CONTROLLED EVIDENCE PERIMETER



Structured data and approved documents feed a cited draft; the model output does not become source evidence.

7. Preserve accountable human decisions

Human review should have defined decision rights. A generic requirement for a human in the loop can become a ceremonial click. The reviewer needs authority, competence, time, source visibility and the ability to reject, correct or stop the output.

The operating matrix assigns preparer, evidence owner, specialist reviewer, compliance or legal reviewer and release authority. A DDQ section about investment process can require an investment owner. A cybersecurity answer can require the technology or security owner. A side-letter interpretation remains with authorised legal or compliance reviewers. A capital notice remains within the finance and administrator approval chain.

The interface should surface the changes and risks that matter: new wording, stale evidence, conflicted sources, unsupported claims, changed numbers, expired approvals and cross-fund leakage. Requiring a reviewer to reread hundreds of unchanged fields weakens attention.

Segregation of duties should survive automation. The person who configures retrieval should not automatically approve every answer. The person who owns a calculation should not bypass independent reconciliation. A vendor administrator should not gain investor-release authority. Smaller teams can use compensating controls such as dual approval, periodic independent sampling and board-level exception review.

Review service levels also need realism. An urgent investor deadline can encourage approval without evidence. The workflow measures queue capacity and routes high-consequence questions early. It can release completed sections while holding unresolved sections when the investor and process permit.

The approval record stores person, role, date, source package, edits, exceptions and release destination. Delegation and absence rules are documented. Emergency processing does not remove retention or later review.

Table 3. Human review matrix

Output	Evidence owner	Required reviewer	Release authority	Stop condition
firm DDQ fact	functional owner	investor operations	authorised IR lead	stale or conflicting source
fund term	legal or fund secretary	legal or compliance	authorised signatory	ambiguous or superseded term
performance value	finance or administrator	performance reviewer	authorised reporting owner	unreconciled calculation
capital notice	administrator and finance	dual financial review	authorised signatory	cash or account mismatch
incident response	incident owner	legal and compliance	designated executive	recipient or obligation uncertain
public claim	evidence owner	compliance and legal	communications approver	unsupported or misleading impression

Roles and segregation should match the manager's organisation and obligations.

8. Control calculations separately from language generation

LP operations contain high-consequence arithmetic: commitments, contributions, distributions, recallable amounts, fees, expenses, ownership, net asset value, performance and allocation. These calculations should remain in deterministic, version-controlled systems with reconciliations and approvals.

AI can help map an incoming template to verified fields, identify variance, compare totals and explain a result. It should not silently create a missing value or modify a formula. The mapping layer stores source field, target field, definition, unit, period, transformation and reviewer.

ILPA's updated Capital Call and Distribution Template seeks greater uniformity and transparency in transaction reporting and includes clearer definitions and streamlined formulas. Its stated implementation timing differs according to adoption of the Performance Template. A manager should map its own systems and fund documents to the applicable template rather than assume the template changes legal economics.

Automated variance checks use explicit thresholds and dependencies. A failed reconciliation holds the affected output and creates an issue. Overrides require reason and approval.

Contractual logic receives a separate rules layer. Most-favoured-nation elections, excuse rights, reporting deadlines, notice periods, confidentiality, advisory committee rights and fee provisions can vary by investor and fund. The obligation register connects each requirement to source clause, trigger, owner, due date and evidence of performance. AI can identify possible relevance; authorised reviewers determine application.

The released response should distinguish a general fund practice from an investor-specific right. Copying a side-letter answer into a general DDQ can expose confidential terms or imply equal treatment where the agreements differ.

9. Protect confidentiality and personal data

LP requests can contain investor identities, commitment amounts, bank details, tax information, beneficial ownership, portfolio company information, employee data and commercially sensitive fund records. Data classification determines whether an AI service can process the information and under which configuration.

The manager records purpose, lawful basis where relevant, data categories, location, retention, access, subprocessors, transfer arrangements, training-use settings, deletion, incident response and data-subject considerations. The ICO's AI and data-protection toolkit offers practical risk support and is under review following the United Kingdom Data Use and Access Act; current guidance should be checked at implementation.


Prompts and outputs are records when they contain business information subject to the applicable retention policy. Sensitive material should not enter personal accounts, consumer tools or unapproved browser extensions. Access follows least privilege and fund segregation.

The security design includes encryption, identity controls, multifactor authentication, logging, secret management, malware protection, prompt-injection defences, export controls, vendor monitoring and tested incident procedures. Retrieved documents are treated as untrusted content until parsed and validated.

Investor-specific access should be enforced upstream of generation. The model should not receive documents that the requesting user and workflow cannot access. Output filtering is a secondary control. Test identities should confirm that an operations analyst, investment professional, external administrator and vendor support user see only their approved perimeters.

Data minimisation applies to both prompts and retained logs. A workflow answering a policy question may not need names, bank details or commitment amounts. Redaction and tokenisation can reduce exposure when they preserve the decision use. Any re-identification path remains controlled.

Figure 4. Confidentiality control zones
THE NARROWEST ZONE INHERITS EVERY OUTER CONTROL



Data movement depends on classification, permission, contract and applicable law.

10. Meet recordkeeping and production requirements

An automated workflow should retain enough information to reconstruct what was received, which evidence was used, which model and configuration acted, what draft was produced, what changed, who approved it and what was released. Logs should be readable and producible throughout the required retention period.

SEC Rule 204-2 materials require registered investment advisers within scope to maintain true, accurate and current books and records, including defined communications and performance records. SEC examination materials ask whether electronic information can be produced promptly and remains readable through software changes. Applicability and retention depend on the adviser and record.

The records model preserves original request, attachments, classification, retrieved sources, prompt template, model identifier, tool actions, generated draft, reviewer edits, approvals, final response and delivery evidence. It avoids retaining unnecessary confidential intermediate material when the policy permits deletion.

Legal hold, investigation and correction processes must suspend ordinary deletion. Vendor exit plans include export of logs, configurations, prompts, evaluation results and answer components in usable formats.

Table 4. AI operations record

Record	Purpose	Minimum lineage	Control owner
incoming request	preserve scope and deadline	sender, time, fund and attachments	investor operations
retrieval package	show evidence used	source, version, passage and access	data owner
model event	reproduce automation	model, configuration, prompt and tools	AI operations
draft and edits	show human judgement	generated version and tracked changes	response owner
calculation link	preserve numerical authority	system output and reconciliation	finance
approval	establish accountability	approver, role, time and exceptions	designated authority
release	evidence delivery	final file, recipient, channel and time	records owner

Retention periods and required records depend on entity, jurisdiction, communication and applicable rules.

11. Test the system as an operational control

Testing begins with a representative request set. It includes common questions, rare questions, ambiguous wording, multiple funds, stale sources, conflicting evidence, tables, calculations, confidential fields, malicious document instructions and missing data.

The evaluation measures retrieval relevance, citation support, fund and period accuracy, field mapping, unsupported statement rate, omission, confidentiality leakage, calculation preservation, abstention, reviewer effort and time. A fluent answer receives no credit when its evidence or boundary is wrong.

Red-team tests attempt cross-fund disclosure, retrieval of unauthorised documents, prompt injection, override of a legal hold, fabrication of a source, use of a superseded answer and automated release without approval. Security and operational teams record results and remediation.

Acceptance thresholds vary by consequence. A formatting assistant can tolerate a different failure profile from a capital-notice control. High-consequence use cases require deterministic controls and can keep generative AI outside the decision path.

12. Monitor drift and exceptions


Models, retrieval indexes, source documents, prompts, integrations and operating teams change. Monitoring should distinguish model drift, data drift, source drift and process drift. A stable model can still produce wrong answers after an outdated document becomes prominent.

Operational measures include request throughput, queue age, evidence gaps, abstention, citations, reviewer changes, corrections, incidents and release failures. Quality sampling examines both accepted and rejected drafts. A low rejection rate can mean strong performance or weak review.

Every exception receives severity, affected fund and investors, owner, due date, release consequence and root cause. Material errors can require suspension, correction, recipient

notification, restatement or specialist review. The incident plan maps privacy, cybersecurity, legal, compliance and investor communication roles.

The change-control board approves material model, prompt, retrieval, data-source and workflow changes. It can require parallel testing before production.



Values are management assumptions used to illustrate monitoring; they are not actual operating results.

13. Govern vendors and concentration risk

The vendor review covers model provider, hosting, retrieval service, vector database, identity, document parsing, monitoring and integration. The manager maps which vendors can access confidential content and which subcontractors support them.

Contract review addresses data use, training, location, access, security, incident notice, audit rights, service levels, intellectual property, confidentiality, deletion, regulatory cooperation, continuity, termination and export. Technical settings should be verified against the contract.

The Financial Stability Board has highlighted third-party dependencies, concentration, cyber risk, model risk, data quality and governance among AI-related financial-sector vulnerabilities. A smaller GP can reduce operational concentration through exportable records, model abstraction, tested manual fallback and alternative providers for critical functions.

Exit testing should confirm that the answer library, evidence metadata, logs, evaluations and workflow definitions can be recovered in usable formats. A vendor failure should slow the process without erasing investor records or approval evidence.

Table 5. Third-party AI control matrix

Control area	Evidence	Operating test	Exit requirement
data use	contract and configuration	verify training and retention settings	deletion certificate
access	roles, keys and logs	periodic privileged-access review	revoke all credentials
resilience	architecture and service levels	outage and manual fallback exercise	portable workflow
security	assurance and incident terms	notification and response simulation	preserve incident records
model change	release and version process	regression evaluation	supported export
subprocessor	current dependency list	change-notice review	data-location confirmation
records	log and archive specification	production and readability test	complete usable export

Contract, privacy and regulatory review should reflect actual services and jurisdictions.

14. Establish AI literacy and operating roles

Staff need training matched to their role. Request coordinators should understand data classification, approved uses and escalation. Reviewers need citation, boundary, unsupported-claim and automation-bias training. Developers need secure design, testing and record controls. Senior managers need consequence, accountability, vendor and incident oversight.

Article 4 of the EU AI Act includes an AI literacy obligation for providers and deployers within its scope, taking account of staff knowledge, experience, education, training and use context. Other frameworks can impose governance and competence expectations through existing rules. The FCA has stated that it relies on existing frameworks, including accountability and governance, for AI in financial services.

Training should include practical failure cases. A reviewer receives a polished draft with the wrong fund, a correct number from the wrong period, an invented citation, an expired side-letter answer and a confidential cross-investor disclosure. Competence is demonstrated through decisions and escalation.

The role register assigns business owner, system owner, data owner, model-risk reviewer, privacy and security contacts, compliance and legal reviewers, records owner and release authorities. The board retains oversight of risk appetite and material use cases.

15. Run a ninety-day implementation

The first thirty days inventory requests, sources, calculations, systems, vendors, jurisdictions and obligations. The manager chooses a low-to-medium consequence pilot, approves use-case boundaries, classifies data and creates a baseline from manual performance.

Days thirty-one to sixty build the request register, evidence metadata, approved-answer library, retrieval filters, citation package, workflow controls and evaluation set. The team tests common, adverse and cross-fund cases in a segregated environment.

Days sixty-one to ninety run a controlled pilot. Every output receives human review. The team compares time, completeness, citations, corrections and exceptions with the baseline. It completes privacy, security, vendor, records and business-continuity reviews before any production approval.


Acceptance requires traceable evidence, accurate boundaries, controlled calculations, reproducible logs, effective stop conditions and named accountability. A faster drafting time alone is insufficient.

The production decision should record intended users, request types, excluded topics, approved data classes, model and retrieval versions, evaluation results, residual risks, monitoring thresholds, fallback and expiry. Approval can be time-limited so material changes trigger reassessment.

The pilot includes a record-production exercise. A reviewer selects one released answer and reconstructs request, evidence, draft, edits, approval and delivery. It also includes a recovery exercise in which the AI service is unavailable and the team completes the request through the manual process.

Figure 6. Ninety-day controlled implementation

PRODUCTION FOLLOWS EVIDENCE AND CONTROL ACCEPTANCE



Timing is a management assumption for planning; actual implementation depends on systems, evidence and obligations.

16. Operate a recurring AI LP-operations office

The operating office triages incoming requests, maintains the answer library, monitors evidence expiry, runs evaluations, manages exceptions and reports material use cases. Daily processing follows queue, due date and consequence. Weekly review examines stale evidence, unsupported fields and reviewer overrides.

Monthly governance reviews quality, security, privacy, vendor changes, incidents and control operation. Quarterly review reassesses use-case classification, tests fallback, reviews access and evaluates changed models and prompts. Annual review refreshes policies, training, risk assessment and board approval.

Event-driven review applies after a fund launch, side-letter change, acquisition, disposal, valuation correction, incident, regulatory development, model change or vendor change. The

system automatically holds affected answer components until review.

The office measures operational outcomes and control quality. It does not claim that automation caused investor commitments without evidence. Investor progression, diligence completion, signed commitments, funded capital and fee collection require separate commercial records.

Capacity planning remains important after launch. Faster intake can create a larger review bottleneck. The office tracks work by consequence, specialist dependency and due date, then adjusts reviewer availability. Automation benefits are realised only when the complete path from intake to approved release improves.

The answer library also needs an owner for retirement. Old funds, expired policies, superseded templates and former team members should be removed from active retrieval while retained according to the records policy. The active corpus and archive serve different purposes.

Table 6. Recurring AI LP-operations office

Cadence	Inputs	Controlled output	Decision owner
daily	requests, evidence and deadlines	routed queue and cited drafts	investor operations
weekly	exceptions, edits and stale answers	remediation and approval pack	functional owners
monthly	quality, privacy, security and vendors	control dashboard	AI governance owner
quarterly	model, prompt, access and fallback tests	continued-use decision	risk committee
annual	policy, training and use-case inventory	board certification and plan	board delegate
event-driven	fund, rule, incident or model change	hold, correction or reapproval	designated authority

Cadence and ownership should reflect the manager's funds, systems and requirements.

17. Use a ten-day diagnostic

The diagnostic collects recent DDQs, investor emails, reporting packs, capital notices, side-letter registers, data-room indexes, answer libraries, policies, vendor contracts and workflow logs. It maps volume, cycle time, owners, repeated questions and failure points.

The team selects a representative request sample and traces each answer to source. It tests fund and period boundaries, calculation authority, confidentiality, record retention, approval and delivery evidence. It also identifies shadow AI use and unapproved data movement.

The output is a board-ready use-case map, consequence tiers, source architecture, priority control gaps, pilot scope, ninety-day plan and recurring ownership. It distinguishes existing capability from proposed capability.

18. Measure value without losing accountability

Operational measures can include median response time, first-pass completeness, percentage of claims with validated citations, reviewer edit rate, exception closure, stale-answer rate, cross-fund errors, rework and on-time delivery. Control measures include access exceptions, prohibited data events, failed evaluations, release blocks and record-production tests.

Cost measures should include model, hosting, integration, security, data preparation, review, vendor management and change control. Saved drafting time can be redeployed to investor dialogue and analysis when the evidence supports that conclusion.

Commercial attribution remains separate. An efficient DDQ process can support a fundraising journey, but an allocation depends on strategy, team, performance, terms, portfolio, risk, market and investor decisions. The manager should record qualified opportunities, diligence progression, signed commitments, funded capital and fees through controlled commercial systems.

19. Common failure modes

The first failure is placing confidential material in an unapproved tool. The second is using the model as a source. The third is retrieving the right concept for the wrong fund. The fourth is allowing a language model to calculate high-consequence values. The fifth is relying on a ceremonial human approval.

Other failures include stale answers, expired documents, invented citations, copied side-letter interpretations, missing prompts and logs, unrestricted vendor training, model changes without regression testing, inaccessible records, no manual fallback and automation of release authority.

The remedy is a bounded operating chain with explicit source, consequence, review, record and stop controls.

20. Board agenda

The board should approve the use-case inventory, prohibited uses, consequence taxonomy, data classifications, source hierarchy, human decision rights, vendor standards, testing thresholds, incident plan, records policy and recurring governance. It should identify the executive accountable for each material workflow.

The immediate questions are practical. Which LP operations create the greatest controlled-value opportunity? Which evidence is authoritative? Which calculations must remain deterministic? Which data cannot enter the proposed service? Which reviewer has genuine authority to stop the output? Which records demonstrate the complete decision chain?

AI can make LP operations faster and more consistent when accountability remains explicit. The durable asset is an evidence and workflow architecture that survives changes in models, vendors, staff and investor templates.

References

- [1] - National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework 1.0. <https://www.nist.gov/itl/ai-risk-management-framework/ai-risk-management-framework-resources>
- [2] - National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, 2024. <https://www.nist.gov/publications/artificial-intelligence-risk->

- [3] - European Union. Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [4] - Financial Conduct Authority. AI and the FCA: Our Approach. <https://www.fca.org.uk/firms/innovation/ai-approach>
- [5] - Financial Conduct Authority. AI in Financial Services: Shaping Our Approach through Industry Engagement, June 2026. <https://www.fca.org.uk/news/blogs/ai-financial-services-approach>
- [6] - Information Commissioner's Office. AI and Data Protection Risk Toolkit. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/ai-and-data-protection-risk-toolkit/>
- [7] - United States Securities and Exchange Commission. Rule 204-2 Books and Records to be Maintained by Investment Advisers. <https://www.sec.gov/files/investment/pfa-vacatur-reverted-rule-text.pdf>
- [8] - United States Securities and Exchange Commission. Questions Advisers Should Ask While Establishing or Reviewing Compliance Programs. https://www.sec.gov/exams/adviser_compliance_questions
- [9] - United States Securities and Exchange Commission. Withdrawal of Predictive Data Analytics Proposed Rulemaking, June 2025. <https://www.sec.gov/rules-regulations/2025/06/s7-12-23>
- [10] - Financial Stability Board. The Financial Stability Implications of Artificial Intelligence, 2024. <https://www.fsb.org/2024/11/fsb-assesses-the-financial-stability-implications-of-artificial-intelligence/>
- [11] - Financial Stability Board. Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector, 2025. <https://www.fsb.org/2025/10/monitoring-adoption-of-artificial-intelligence-and-related-vulnerabilities-in-the-financial-sector/>
- [12] - Financial Stability Board. Sound Practices for Financial Institutions' Responsible AI Adoption: Consultation Report, 2026. <https://www.fsb.org/uploads/P100626.pdf>
- [13] - Institutional Limited Partners Association. Due Diligence Questionnaire 2.0. <https://ilpa.org/resources-tools/resource-library/due-diligence-questionnaire/>
- [14] - Institutional Limited Partners Association. Updated Capital Call and Distribution Template, 2025. <https://ilpa.org/industry-guidance/templates-standards-model-documents/ilpa-templates-hub/ilpa-capital-call-distribution-template/>
- [15] - Institutional Limited Partners Association. Reporting Template. <https://ilpa.org/reporting-template/>

About the Author

Chennakeshav Adya, Independent Researcher