MatchPoint Research Working Paper No. 2026/T13

AI for Fund Marketing and LP Relations: Personalisation at Scale

A controlled evidence, recipient and release architecture for fund managers and family-office allocators

Chennakeshav (CK) Adya
Independent Researcher
August 2026

ABSTRACT

Background. Fund managers communicate strategy, performance, risk and terms to limited partners with different mandates, jurisdictions and diligence priorities. **Objective.** This paper develops a controlled method for AI-assisted fund marketing and LP relations at scale. **Approach.** The analysis reviews 32 primary regulatory, official industry, AI-governance and task-specific research sources available through 1 August 2026 and translates them into an evidence-to-release operating model. **Findings.** Personalisation should begin with approved evidence and recipient eligibility. Each material claim inherits source conditions, deterministic checks and named review. The accountable productivity unit is a quality-adjusted accepted LP communication packet. **Implications.** GCC fund managers can begin with meeting preparation, DDQ retrieval and controlled follow-up. Family-office CIOs can use the same architecture to organise manager evidence and expose gaps. Attributed revenue, cost reduction and loss reduction remain USD 0 until approved observed attribution exists.

Highlights

Unit. The accountable unit is a quality-adjusted accepted LP communication packet.

Sequence. Evidence, eligibility, audience, channel, approval and retained record form one controlled workflow.

Claims. Material statements inherit source version, scope, caveat, permission and expiry.

Personalisation. Depth rises only with verified recipient facts, current evidence and review capacity.

Economics. Revenue, cost reduction and loss reduction remain zero until observed attribution is approved.

JEL Classification: G23, G24, M31, O33, C81, L86

Keywords: fund marketing, limited partners, investor relations, personalisation, generative AI, due diligence, financial promotions, data governance, family offices, fund placement

1. INTRODUCTION

Fund marketing and limited-partner relations are evidence-intensive operating systems. A fund manager must explain the strategy, team, governance, portfolio, performance, terms, risks and reporting arrangements to prospective and existing investors whose mandates, knowledge, legal status and diligence priorities differ. The same manager may answer a family-office chief investment officer, an institutional consultant, a private bank platform and a sovereign allocator in the same week. Relevance therefore matters. Consistency, substantiation and authorised distribution matter at least as much.

Generative artificial intelligence can assemble prior answers, identify relevant approved passages, prepare meeting briefs, classify requests, propose follow-up and maintain relationship records. It can also create unsupported performance language, lose the conditions attached to a number, misclassify a recipient, reuse confidential material or scale a weak control across hundreds of communications. The operational question is how to obtain useful personalisation while preserving evidence, eligibility, approval and a retained record.

This paper develops a controlled framework for **AI for Fund Marketing and LP Relations: Personalisation at Scale**. The primary ideal customer profile is B2, GCC fund managers and general partners raising private-equity, private-credit or venture-capital funds. The secondary profile is A2, family-office chief investment officers and heads of alternatives assessing funds and direct opportunities. The framework is intended for professional and institutional contexts. It does not determine whether a communication, recipient, fund or activity falls within a legal exemption or regulatory perimeter.

The central unit of work is a **quality-adjusted accepted LP communication packet**. A packet contains the request or purpose, recipient classification, approved evidence, generated draft, mandatory disclosures, reviewer decision, released version and retained record. Counting words, prompts, drafts or messages rewards output volume. Counting accepted packets connects productivity to material that a named human reviewed, released and can reproduce.

Five propositions organise the analysis. First, personalisation should begin with evidence eligibility and recipient eligibility, not free-form generation. Second, every claim should inherit the provenance, date, scope, caveats and permissions of its source. Third, the personalisation depth should rise only when evidence quality, recipient certainty and review capacity support it. Fourth, customer-facing release remains a named human decision. Fifth, realised economic value requires observed and approved attribution. This paper therefore records attributed revenue, attributed cost reduction and attributed loss reduction as USD 0.

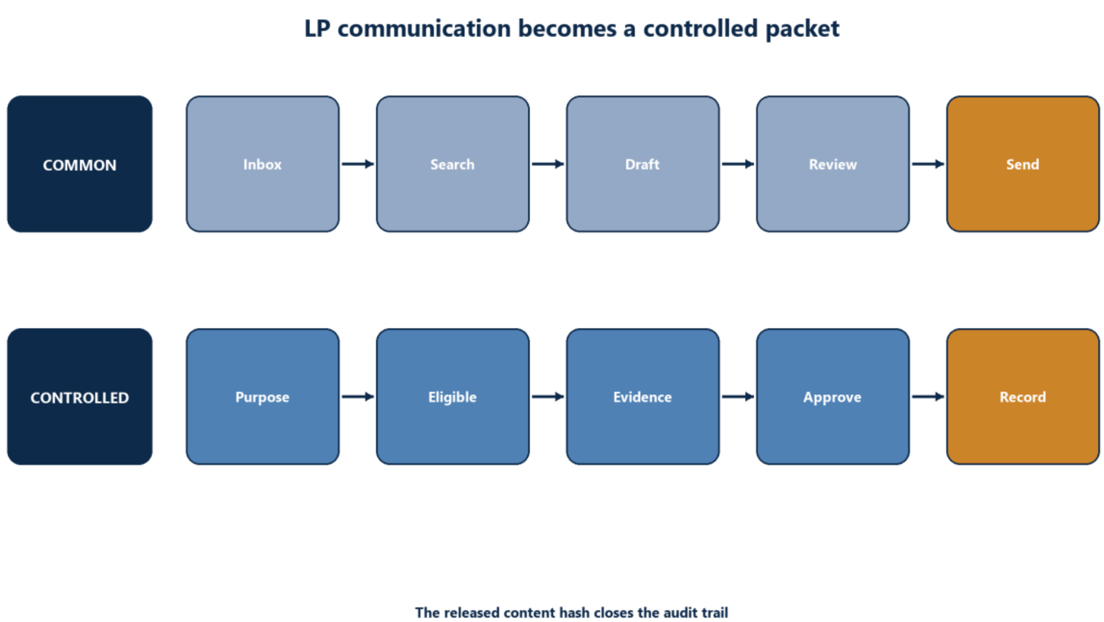

Figure 1. Before and after LP communication workflow
Source: Matchpoint framework.

The analysis reviews primary regulatory material, official industry standards, official privacy guidance, recognised AI-governance frameworks and task-specific productivity studies available through 1 August 2026. Regulatory requirements vary by entity, activity, jurisdiction, recipient and channel. The paper preserves those boundaries. It offers an operating model for legal and compliance assessment rather than a legal conclusion.

The contribution is practical. Sections 2 and 3 define the evidence boundary and LP journey. Section 4 maps use cases and stop conditions. Section 5 specifies a controlled architecture. Section 6 introduces a personalisation ladder. Section 7 defines measurement and illustrative economics. Section 8 provides B2 and A2 playbooks. Section 9 consolidates legal, regulatory and data-governance questions. Section 10 presents a gated adoption roadmap. Sections 11 and 12 state limitations and conclusions. Appendices provide reusable registers, packet fields and scorecards.

2. SCOPE, DEFINITIONS AND EVIDENCE BOUNDARIES

2.1 What this paper means by fund marketing and LP relations

Fund marketing covers the preparation, approval and distribution of communications intended to explain or promote a fund, strategy or manager to prospective investors. LP relations covers the continuing exchange of information with existing investors, including onboarding, data-room access, capital activity, portfolio reporting, meetings, questions and periodic communications. The boundary between education, relationship management, marketing, solicitation and regulated advice depends on facts and applicable law. The workflow must retain the intended purpose and recipient because the same sentence can have different implications in different contexts.

Personalisation means selecting, ordering and explaining approved information for a defined recipient and purpose. It includes mandate-aware topic selection, jurisdiction-appropriate disclosures, role-specific terminology, meeting context and the recipient's recorded questions. It excludes invented facts, hidden persuasion, unapproved performance transformations, false urgency and the unsupported inference of sensitive characteristics.

Artificial intelligence in this paper includes retrieval, classification, extraction, summarisation and generation systems. A language model may propose prose. Deterministic services should perform arithmetic, version comparison, eligibility rules, required-field checks and other tasks where a fixed rule is available. Human owners approve policy, source material, recipient status, material communications and release.

Term	Operational definition	Evidence retained	Named authority
Approved evidence object	Versioned source passage, table, metric or answer cleared for a stated use	Source, owner, date, scope, caveat, permission, expiry	Content owner and compliance
Recipient profile	Recorded institutional facts relevant to communication eligibility and relevance	Entity, role, jurisdiction, professional status, mandate, consent and source	Relationship owner and compliance
Communication packet	Complete work item from purpose through retained released version	Request, sources, draft, checks, reviewer, release and archive	Named reviewer
Personalisation rule	Approved mapping from recipient fact to content selection or format	Rule version, rationale, inputs, exclusions and test	Product owner and compliance
Material claim	Statement whose error could affect investment, legal, reputational or relationship decisions	Exact source, calculation, conditions and reviewer	Claim owner
Release	Authorised external or internal delivery of a versioned communication	Recipient, channel, time, file hash and authority	Named releaser

2.2 Regulatory and standards map

The United States Securities and Exchange Commission marketing rule applies to advisers registered or required to register under the Advisers Act. The SEC small-entity guide explains that advertisements can include communications offering advisory services to prospective clients or private-fund investors, subject to definitions and exclusions [1]. The rule contains general prohibitions and conditions for testimonials, endorsements, third-party ratings and performance information. SEC staff frequently asked questions provide current staff views, including gross and net performance presentation questions; the FAQ expressly states that it has no legal force or effect [2].

FINRA Rule 2210 separates correspondence, retail communications and institutional communications and requires communications to be fair and balanced and not misleading within its scope [6]. Applicability depends on the communicating entity, channel and recipient. A fund manager should not apply a single global template merely because an LP is sophisticated.

In the United Kingdom, the FCA's social-media guidance states that financial promotions must remain fair, clear and not misleading and that the rules are technology-neutral [7]. COBS 4 contains requirements concerning communication with clients and financial promotions, including fair presentation, risk prominence and suitability of presentation for the likely audience [8]. COBS 4.10 addresses systems and controls for approving and communicating promotions [9]. These sources require case-specific application by qualified advisers.

EU Regulation 2019/1156 requires relevant fund marketing communications to be identifiable as such, to describe risks and rewards with equal prominence and to be fair, clear and not misleading [10]. ESMA guidance expands on presentation and consistency expectations [11]. In the Dubai International Financial Centre, the current DFSA Conduct of Business module includes required firm identity and professional-client statements for marketing material and reasonable steps concerning who receives professional-client material [12]. The DFSA collective-investment-law provisions may also be relevant to offers and fund marketing [13].

The Institutional Limited Partners Association Due Diligence Questionnaire 2.0 standardises many diligence enquiries and expressly leaves investment determination with the LP [21]. ILPA's updated performance and reporting templates create structured fields that can reduce arbitrary restatement while preserving the manager's responsibility for accurate reporting [22-24]. These are industry standards rather than legislation.

Source family	Verified contribution	Boundary carried into this paper
SEC Marketing Rule and guide [1-5]	General prohibitions, performance, testimonials, endorsements, ratings, recordkeeping and observed examination themes	Adviser, communication and recipient scope must be assessed; staff FAQs have no legal force
FINRA Rule 2210 [6]	Communication categories and fair, balanced, non-misleading standard	Applies within FINRA member scope and defined categories
FCA and COBS [7-9]	Technology-neutral financial-promotion controls, audience, risk and approval expectations	UK perimeter and exemption analysis remains external legal work
EU and ESMA [10,11]	Identifiability, risk-reward prominence, consistency and presentation	Jurisdiction, fund type and cross-border activity matter
DFSA materials [12,13]	Marketing-material statements, recipient controls and fund-offer context in the DIFC	DIFC scope cannot be extended to the rest of the UAE without analysis
Privacy authorities [14-20]	Purpose, transparency, objections, electronic marketing and profiling governance	Applicable controller, subscriber type, jurisdiction and lawful basis require assessment
ILPA frameworks [21-24]	Standard diligence, performance, reporting and portfolio-company fields	Voluntary industry standards; manager and LP judgement remain necessary
NIST and IOSCO [25-27]	Risk-management and capital-markets AI governance considerations	Frameworks and consultation material do not establish compliance
Task studies [28-30]	Productivity evidence in selected writing, support and consulting tasks	No direct evidence of fund-raising conversion, LP satisfaction or realised revenue

2.3 Evidence hierarchy and claim discipline

An operating team should rank evidence before a model uses it. Executed fund documents, filed or approved regulatory materials, audited financial information and formally approved investor reports sit above drafts, meeting notes and public commentary. Each evidence object should carry a jurisdiction, period, entity, fund, share class, calculation method, approval state, confidentiality class, recipient permissions and expiry.

The system should refuse to promote a draft into an approved source. It should also refuse to combine numbers whose definitions differ. Gross performance and net performance, realised and unrealised value, fund-level and deal-level return, target and actual outcome, and preliminary and audited figures are distinct objects. A citation to the correct document is insufficient when the extracted value has lost its unit, period, currency, denominator or condition.

Evidence tier	Example	Permitted initial use	Release requirement
E0 prohibited or unknown	Unattributed claim, expired file, unclear owner	Quarantine and investigate	Cannot be released
E1 working input	Meeting note, analyst draft, unverified CRM field	Internal triage with visible label	Independent verification
E2 controlled reference	Approved policy, signed-off standard answer, current legal wording	Drafting within stated purpose	Current recipient and channel check
E3 authoritative fund evidence	Executed documents, audited report, approved performance source	Material claim drafting	Exact conditions, reviewer and retained citation
E4 external authority	Current law, regulator page, official standard	Control design and disclosure context	Applicability assessment and source date

The evidence cut-off for this paper is 1 August 2026. Later changes are outside scope. The author identity is unverified pending CK approval. No client, fund or LP dataset was supplied for this research. All worked economic scenarios are unverified illustrative management assumptions.

3. THE LP JOURNEY AND CONTROLLED DATA MODEL

3.1 A journey composed of decisions

The LP journey is often represented as a funnel: target, contact, meeting, diligence, investment committee and close. That view omits the control decisions that determine whether a communication should exist. A controlled journey records at least seven gates: evidence readiness, recipient eligibility, contact permission, purpose, content selection, human approval and archive. Relationship progression is one outcome. An accountable record is required at every outcome, including decline, defer and no response.

At the earliest stage, a manager may build a market map from permitted institutional sources. The system should distinguish a public business address from a personal address and a corporate subscriber from an individual or sole trader where relevant. It should record the provenance and date of a contact. It should avoid inferring wealth, health, ethnicity, religion, politics or other sensitive attributes. A role such as chief investment officer is a professional fact when sourced appropriately; an inferred preference built from private behaviour can carry a different privacy and ethical risk.

During qualification, the relationship owner records stated mandate features: asset classes, geography, ticket range, liquidity, exclusions, reporting expectations and decision process. These statements should carry their source and date. Silence is not evidence. A model may propose a question to close a gap. It should not convert a missing field into an asserted preference.

During diligence, the packet draws from the data room, current fund documents, approved DDQ responses and controlled performance sources. Every response should identify whether it answers the question completely, partially or through an attached source. A repeated answer should inherit the current version and caveats. When an LP asks for a new calculation or side-letter commitment, the system routes the request to the relevant owner and blocks generation of an agreement.

After commitment, LP relations continues through onboarding, notices, reporting, meetings and ad hoc questions. Personalisation can improve navigation and explanation. Equal-treatment duties, side-letter terms, most-favoured-nation processes, confidentiality and selective-disclosure questions require legal and operational controls outside the model.

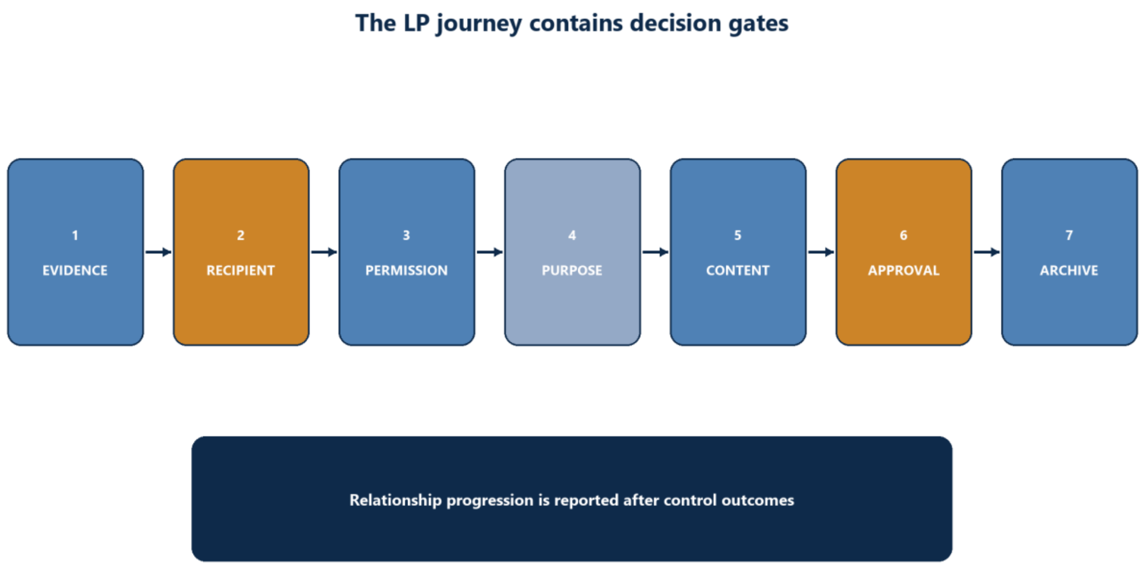

Figure 2. LP journey with evidence, recipient and release gates

Source: Matchpoint framework; regulatory context [1-20].

3.2 The recipient profile

The recipient profile should be small, sourced and purpose-bound. A large behavioural profile creates stale data, access risk and hidden inference. A useful profile begins with entity identity, verified work role, jurisdiction, recipient category, relationship stage, stated mandate, interaction preferences and current open questions. Every field needs a source, owner, timestamp and review date.

Recipient-profile field	Acceptable source example	Use	Stop condition
Entity and work role	Corporate website, signature block, verified CRM entry	Addressing and role-relevant ordering	Identity conflict or stale employment
Jurisdiction and recipient category	Onboarding or compliance record	Eligibility and disclosure routing	Missing or unapproved classification
Mandate	LP-stated request, RFP, meeting note confirmed by owner	Topic selection and question preparation	Inferred mandate presented as fact
Ticket range	LP-stated range with date	Relevance screening	Unverified estimate
Exclusions	Approved mandate note	Avoid irrelevant materials	Ambiguous scope
Preferred channel	Recorded permission or request	Delivery routing	Consent or channel status unclear
Open questions	Meeting record or inbound request	Follow-up composition	Request requires new promise or calculation

Recipient-profile field	Acceptable source example	Use	Stop condition
Confidentiality	Agreement and access-control record	Source and attachment permissions	Source-recipient mismatch

3.3 The approved evidence object

An evidence object is the atomic unit supplied to retrieval. It should include the exact text or value plus metadata. For a performance number, metadata includes fund, vehicle, period, currency, gross or net basis, fee assumptions, valuation date, calculation method, source file and approver. For a team biography, metadata includes role, start date, scope, approved wording and public or confidential status. For a case study, metadata includes permission, anonymisation decision, realised status and required caveats.

Versioning is essential. A current approved DDQ answer may supersede three earlier answers without deleting the audit history. Retrieval should use the current version by default and surface conflicts. The packet should record the object identifiers and versions used. A later source update does not silently rewrite a communication already released.

3.4 The message record

The released communication should be reproducible. The archive records input request, recipient ID, profile version, selected evidence objects, model and prompt version, deterministic checks, draft versions, reviewer edits, disclosures, release channel, delivery time and final content hash. Sensitive prompts and model telemetry require controlled retention and access. The record exists to support supervision, correction, complaint handling and learning.

Packet state	Entry condition	Permitted action	Exit evidence
Requested	Inbound question or approved campaign purpose	Classify request	Purpose and owner
Eligible	Recipient, channel and activity cleared	Select sources	Eligibility record
Grounded	All material claims linked to approved objects	Draft	Evidence manifest
Checked	Deterministic and policy tests pass	Human review	Check results and exceptions
Approved	Named reviewer accepts exact version	Release	Approval identity and timestamp
Released	Approved version delivered through authorised channel	Archive and monitor	Recipient, channel, content hash
Closed	Response, expiry, withdrawal or correction recorded	Report metrics	Outcome and follow-up decision

4. USE-CASE FRONTIER

4.1 Investor mapping and research

AI can classify public or licensed institutional information into an initial market map, identify likely mandate questions and summarise prior permitted interactions. The useful output is a research packet with sources, dates and unknowns. A contact list without provenance or eligibility creates operational risk. The model should state that a mandate fit is unverified until the recipient or an approved source confirms it.

For B2 managers, mapping can reduce repeated research across a fundraising team. For A2 family offices, the same architecture can organise the manager universe, flag missing evidence and prepare comparison questions. Neither use authorises a recommendation. Ranking criteria should be declared. Sponsored access, commercial relationships or placement incentives should be disclosed where relevant.

4.2 Data-room navigation and diligence response

Retrieval can locate the current source for a question and draft an answer with linked passages. ILPA DDQ fields provide a useful taxonomy for organisation, strategy, team, governance, operations, service providers, responsible investment and technology [21]. The system should preserve the LP's wording and indicate complete, partial or unavailable evidence.

Answers involving performance, track record, valuation, fees, conflicts, legal terms, key-person provisions, side letters or regulatory status carry high materiality. They require specialist review and exact source conditions. The model should not reconstruct a missing value from surrounding context. A blank field should remain blank with an owner and due date.

Use case	AI role	Deterministic control	Human authority	Prohibited shortcut
DDQ triage	Classify question and retrieve candidate sources	Required fields, version and permission	Diligence owner	Treat semantic similarity as an answer
RFI response	Draft source-linked response	Number, unit, date and cross-document checks	Content and compliance reviewer	Invent a missing figure
Data-room guide	Explain location and document purpose	Access-control check	Data-room owner	Expose another recipient's material
Performance response	Assemble approved table and wording	Recalculate from controlled source	Finance, compliance and releaser	Transform gross to net through free-form generation
Terms question	Retrieve executed or approved wording	Document-version match	Legal and fund owner	Negotiate or promise terms

4.3 Meeting preparation

A meeting brief can combine the verified recipient profile, prior questions, current fund stage, relevant materials and required disclosures. It should distinguish recorded facts from proposed hypotheses. For example, “the LP stated a preference for lower-duration credit on 14 June” is a sourced fact. “The LP may prefer this strategy” is an inference and should not appear as confirmed context.

The brief should include open contradictions and uncomfortable questions. A productive system does not optimise only for positive narratives. It should surface a changed team member, valuation dispersion, delayed exit, concentration, cash-flow issue, revised target or inconsistency between a deck and data room. The relationship owner decides how to address it with approved facts.

4.4 Follow-up and next-best action

AI can draft a concise follow-up, list agreed actions and route owners. Next-best action should be a constrained workflow recommendation, not autonomous persuasion. The system can propose “send approved portfolio-construction note after compliance review” or “request confirmation of mandate range.” It should not apply false scarcity, infer emotional vulnerability or schedule repeated messages beyond approved contact policy.

The follow-up must reflect the meeting record. If meeting notes are unverified, the draft carries that label until the relationship owner confirms them. Commitments, timelines and data requests become tasks with owners. The system should not silently turn a discussion point into a firm promise.

4.5 Periodic LP communications

Quarterly letters, annual meetings and portfolio updates offer a high-value reuse opportunity because core evidence repeats across recipients. Personalisation can add an approved portfolio-relevance summary, navigation aid or role-specific appendix. The canonical report remains controlled. Any variation should preserve the same underlying facts, risk balance and approval date.

ILPA reporting and performance templates can create structured inputs [22-24]. Their adoption does not validate the data. The workflow should reconcile structured fields to the accounting, administration and valuation sources approved by the manager. Recipient-specific explanations should not alter reported values.

4.6 CRM hygiene and relationship intelligence

Models can extract entities, dates, questions and actions from approved correspondence and propose CRM updates. Every proposed update should be reviewed or subjected to a rule appropriate to its materiality. The CRM should retain source links and distinguish stated facts, analyst assessments and system-generated suggestions. A relationship score without interpretable inputs can conceal bias and stale data.

4.7 Stop conditions

The safest automation includes explicit reasons to stop. A stop does not represent technical failure. It is a designed outcome when evidence or authority is insufficient.

Stop condition	Required response	Owner
Recipient category or jurisdiction unresolved	Hold distribution and obtain compliance classification	Compliance
Source expired, conflicting or unapproved	Quarantine claim and route to content owner	Content owner
Performance basis unclear	Block draft and obtain controlled calculation	Finance and compliance
New promise, term or side-letter request	Route without proposed agreement language	Legal and fund owner
Confidentiality mismatch	Deny retrieval and alert data owner	Data-room owner
Sensitive personal-data inference	Remove inference and assess lawful processing	Privacy owner
Model cannot support exact citation	Produce an evidence-gap response	Research owner
Required human reviewer unavailable	Queue or decline; do not release	Workflow owner

5. CONTROLLED ARCHITECTURE

5.1 Evidence to release

The architecture has seven layers: source systems, approved evidence registry, recipient and permission registry, retrieval and policy orchestration, generation, deterministic validation, and human release. Each layer should have a named owner and observable state. The model should receive the minimum information needed for the packet.

Source systems may include the fund data room, document management system, CRM, fund administration outputs, approved market data and compliance library. The evidence registry does not copy everything indiscriminately. It indexes approved objects and their conditions. The recipient registry supplies verified classification, mandate and permissions. The orchestrator selects sources only after recipient and purpose checks.

Generation produces a proposed draft plus a claim manifest. Deterministic services verify dates, units, repeated figures, required disclosures, recipient-specific restrictions, broken citations and approved versions. A policy engine applies materiality and routes the packet. The reviewer sees the draft, changes and evidence together. Release occurs only for the reviewed content hash.

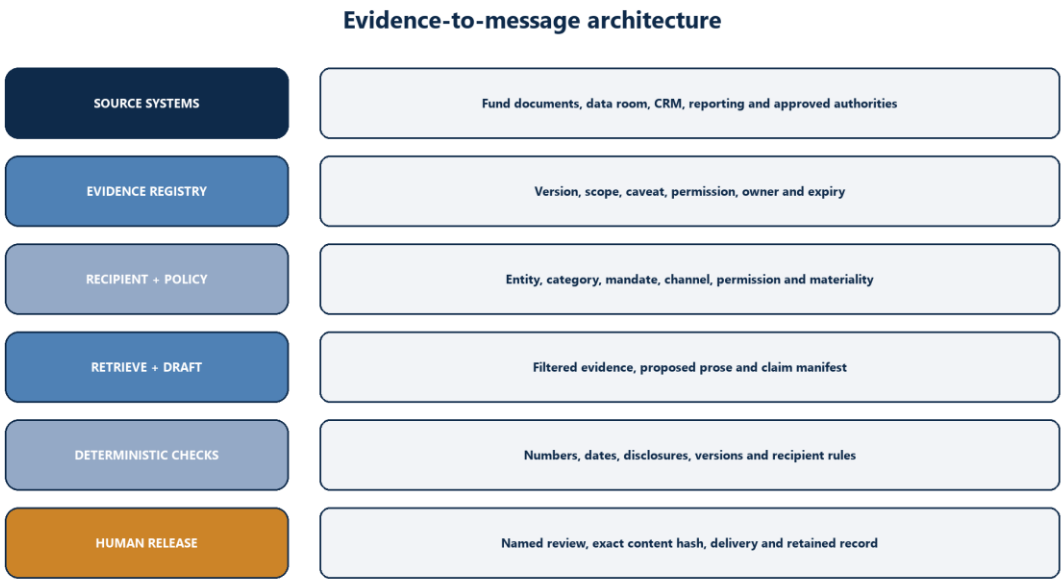

Figure 3. Evidence-to-message architecture

Source: Matchpoint architecture; governance context [25-27].

5.2 Retrieval and claim manifests

Retrieval quality should be measured at the claim level. Topical relevance is one test. Permission, currency, authority and completeness are separate tests. A retrieved passage from a superseded deck can be semantically perfect and operationally wrong. The retriever should filter before ranking, surface source status and return an explicit no-answer result.

A claim manifest lists each material proposition, source object, passage, date, scope, calculation and caveat. It also records generated connecting language that does not need a factual citation. The manifest lets reviewers focus on material changes. It supports later correction when an underlying object changes.

Manifest field	Purpose	Example control
Claim ID and exact text	Stable review unit	Hash exact sentence or table cell
Materiality	Route scrutiny	High for performance, terms, risk and regulatory status
Evidence object/version	Reproducibility	Approved current version required
Passage or calculation	Direct substantiation	Page, cell or deterministic function retained
Scope and caveat	Prevent overgeneralisation	Fund, period, currency, gross/net basis
Recipient permission	Protect confidential content	Entitlement must cover the object
Reviewer decision	Human accountability	Accept, edit, reject or escalate
Expiry and correction path	Ongoing control	Alert owner and identify released packets

5.3 Tool and model boundaries

The language model should draft, classify and explain. It should not serve as the system of record, calculator, permission authority or delivery credential. Tool calls should be typed, limited and logged. Retrieval should be read-only for the drafting agent. CRM updates, data-room grants and message sends should use separate authorised actions with explicit approval.

Model prompts and outputs can contain confidential information. Deployment design should address provider terms, retention, training use, regional processing, subcontractors, encryption, access, incident response and exit. The paper does not assess a specific provider. A model label alone does not establish confidentiality, security or regulatory fitness.

5.4 Validation suite

Validation should combine a reference set with adversarial cases. Reference cases test recurring work such as strategy explanations, approved track-record questions, DDQ fields and meeting follow-up. Adversarial cases test stale sources, conflicting metrics, similar fund names, unauthorised recipients, prompt injection in documents, hidden instructions, performance transformation, legal promises and sensitive inference.

The suite should measure retrieval recall for required evidence, unsupported material-claim rate, citation accuracy, numeric consistency, disclosure completeness, recipient-policy compliance, reviewer acceptance, correction rate and turnaround time. An average score can conceal a severe failure. High-materiality cases need zero-tolerance or separately governed thresholds.

Test family	Metric	Minimum release evidence
Retrieval	Required-source recall and current-version precision	Approved reference answers and failure analysis
Grounding	Unsupported material-claim rate	Claim-level audit with severity
Numbers	Exact match, units, periods and basis	Deterministic reconciliation log
Recipient	Eligibility and permission-policy pass rate	Representative jurisdiction and role cases
Disclosure	Required text and prominence	Channel-rendered inspection
Robustness	Injection, conflict and ambiguity handling	Red-team cases and stop evidence
Human factors	Reviewer acceptance, edit distance and missed-error rate	Named reviewers and calibrated sample
Operations	Latency, queue age, correction and rollback	Monitoring and incident drill

5.5 Access, logging and retention

Least privilege applies to documents, tools and releases. Relationship staff may see the materials needed for assigned recipients. Technical administrators should not receive broad content access by default. Logs should be useful for accountability while avoiding unnecessary replication of sensitive prompts. Retention should follow applicable recordkeeping, privacy, contractual and litigation requirements assessed by qualified owners.

The system needs a kill switch. Owners should be able to stop generation, block a source, revoke recipient access, suspend a channel and identify affected packets. A correction workflow should generate a controlled notice where appropriate, preserve the original record and state who authorised the correction.

6. PERSONALISATION DESIGN

6.1 A controlled ladder

Personalisation should advance through five levels. Level zero provides the canonical approved document. Level one changes navigation, language and format. Level two selects approved modules according to verified mandate. Level three adds a source-linked contextual explanation. Level four prepares a bespoke response using controlled calculations and specialist review. Autonomous promises, legal terms, performance transformations and recommendations sit outside the ladder.

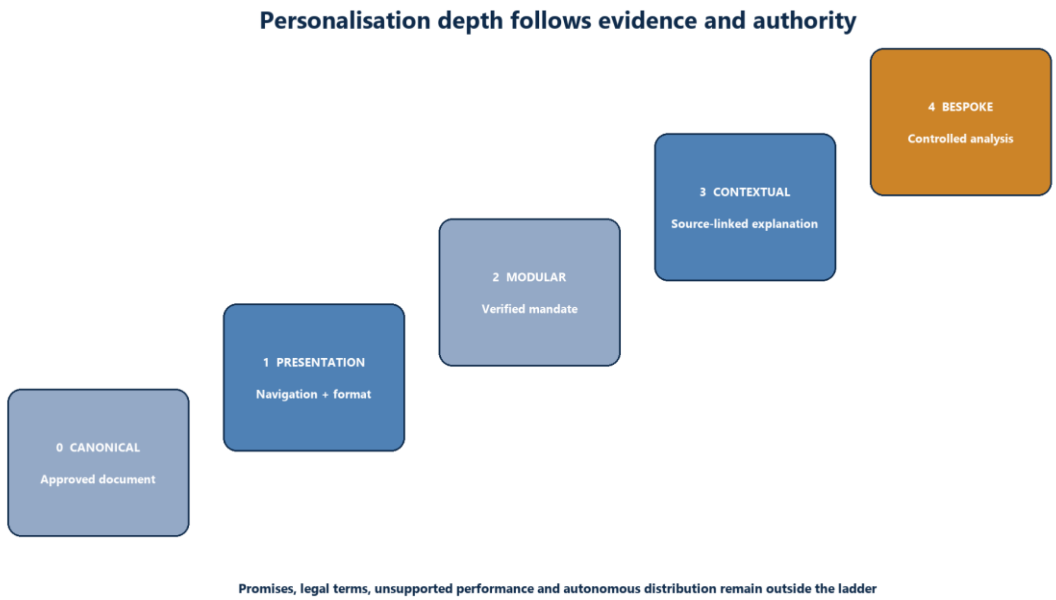

Figure 4. Personalisation control and relevance ladder

Source: Matchpoint framework.

Level	Output	Permitted inputs	Review	Example
0 Canonical	Same approved document	Recipient and permission only	Existing publication approval	Current PPM or standard report
1 Presentation	Navigation, role vocabulary, language or length	Verified role and requested format	Template and spot review	CIO summary plus links to full report
2 Modular	Ordered approved modules	Stated mandate and stage	Relationship-owner review	Credit sections before venture sections
3 Contextual	Recipient-specific explanation tied to sources	Confirmed questions and approved evidence	Content and compliance review	How strategy liquidity aligns with stated policy question
4 Bespoke	New analysis or response	Controlled data, calculations and legal scope	Specialist multi-owner approval	Custom exposure table requested in diligence

The ladder separates relevance from persuasion. A recipient may appreciate a concise route to the correct evidence. The system should avoid exploiting inferred behavioural traits. Personalisation rules should be understandable to the reviewer and testable across recipient groups.

6.2 Content modules

A content module should have an owner, approved wording, required facts, allowed variants, prohibited uses, mandatory risk context and expiry. Examples include strategy summary, team, portfolio construction, sourcing, value creation, risk, track record, fees, operations, responsible investment and reporting. The generator can join modules through approved transitions while preserving each module's conditions.

Long-form and short-form versions need separate approval where material meaning can change. A social post, email, one-page summary and DDQ answer do not become equivalent merely because they cite the same paper. Channel constraints affect prominence, context and who may receive the communication.

6.3 Recipient and mandate relevance

The system should apply only explicit, professionally relevant facts. For a family-office CIO, useful confirmed dimensions may include governance process, liquidity budget, concentration limits, co-investment interest, reporting cadence and risk questions. The manager should obtain these through stated mandates and permitted interaction. Estimated wealth, private family circumstances or inferred sentiment should not guide automated persuasion.

The family office receiving a personalised communication also needs transparency. It should be able to request the canonical source, understand which claims are manager-supplied and identify missing information. A2 teams can apply the same evidence architecture to compare managers without accepting the manager's relevance framing as an investment judgement.

6.4 Channel controls

Email, data rooms, portals, social platforms, messaging applications and meetings have different disclosure, permission, rendering and retention characteristics. The approved channel matrix should state audience, allowed content levels, required disclosures, link handling, attachment restrictions, approval and archive method.

Channel	Suitable controlled use	Principal risk	Control
Data room	Recipient-entitled diligence material	Cross-recipient leakage and stale documents	Entitlement, watermark, version and access log
LP portal	Existing-investor reports and notices	Identity, availability and selective disclosure	Strong authentication, canonical files and incident plan
Email	Approved follow-up and requested material	Wrong recipient, forwarding and electronic-marketing rules	Recipient check, approved sender, suppression and archive
Social platform	Public or permitted audience education	Context loss, audience spread and promotion rules	Channel-specific approval, balanced risk and durable landing page
Messaging application	Agreed operational follow-up	Informality, retention and device risk	Restricted content, approved account and archive
Meeting assistant	Briefing and controlled notes	Recording, consent and inaccurate extraction	Notice, permission, review and deletion policy

6.5 Fairness and distribution quality

Personalisation models can allocate attention unevenly. A score trained on historical conversions may favour familiar institutions, geographies or communication styles. The firm should examine who receives access, response speed, senior attention and complete information. Commercial prioritisation may be legitimate within policy. It should be explicit, monitored and consistent with legal and contractual duties.

Quality sampling should include declined and low-score recipients. A system that evaluates only successful communications cannot identify exclusion or misclassification. Complaints, corrections and opt-outs should feed governance, not become hidden negative labels.

7. MEASUREMENT AND ECONOMICS

7.1 The quality-adjusted unit

The numerator is the number of packets released and accepted by the defined review process. The denominator can be analyst hours, reviewer hours or elapsed business time. A packet loses quality credit when it contains an unsupported material claim, wrong recipient, missing disclosure, incorrect number, confidentiality breach or required correction. The firm should report volume and quality separately.

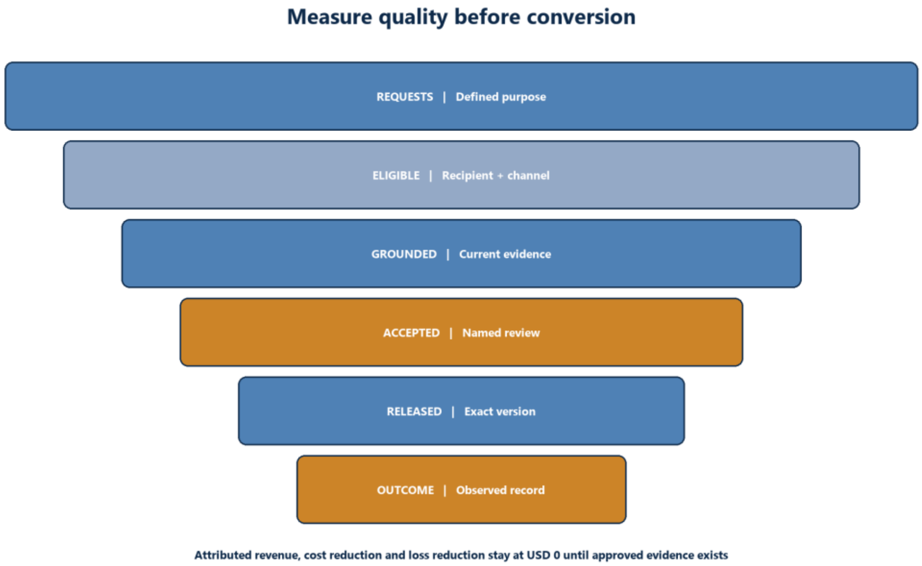

Figure 5. Engagement funnel with quality gates
Source: Matchpoint measurement framework; task evidence [28-30].

Metric	Definition	Evidence required	Misleading proxy to avoid
Accepted packet rate	Reviewer-accepted packets divided by submitted packets	Workflow decisions and versions	Generated drafts
First-pass acceptance	Accepted without material edit	Review diff and materiality policy	Grammar score
Grounding pass	Packets with zero unsupported material claims	Claim audit	Citation count
Numeric pass	Packets with correct values, units, periods and basis	Deterministic checks	Apparent fluency
Turnaround	Request to approved release, excluding agreed waiting states	Event timestamps	Model response time
Reviewer load	Active review minutes per accepted packet	Time or structured activity record	Total staff hours without allocation
Recipient outcome	Requested next step, complete response, opt-out or correction	CRM and communication record	Open rate alone
Attributed economics	Approved revenue or cost outcome causally assigned under policy	Finance-approved attribution record	Pipeline value

7.2 External productivity evidence

Noy and Zhang studied 444 professionals completing selected writing tasks and reported faster completion and higher assessed quality with generative AI [28]. Brynjolfsson, Li and Raymond studied 5,179 customer-support agents and reported an average productivity gain concentrated among less-experienced workers [29]. Dell'Acqua and co-authors studied knowledge workers across tasks inside and outside an AI capability frontier and found that performance effects varied by task [30]. These studies support testing task-specific productivity and heterogeneity. They do not demonstrate fund-marketing conversion, LP trust, investment performance or realised revenue.

The firm should therefore establish a local baseline before adoption. It should sample comparable task types and materiality, record manual and assisted cycle time, review quality blind where practical, and observe corrections after release. Selection bias matters: teams may choose easier cases for the pilot. Learning effects and reviewer behaviour also affect results.

7.3 Illustrative management scenario

The following scenario is unverified and illustrative. It is a design aid for measurement, not a forecast. Assume a fundraising team processes 120 packets per month. A manual packet uses 2.4 analyst hours and 0.6 reviewer hours. An assisted packet uses 1.4 analyst hours and 0.8 reviewer hours because evidence review becomes more explicit. Assume 70 per cent of packets become eligible for the assisted process after exclusions. The scenario implies gross analyst time released and additional review demand. It does not establish cash savings because staff may redeploy time and because implementation, validation and governance costs are unknown.

Illustrative input	Unverified assumption	Calculation use
Monthly packets	120	Workload scale
Eligible share	70%	84 assisted packets
Manual analyst time	2.4 hours	Baseline comparison
Assisted analyst time	1.4 hours	Scenario comparison
Manual reviewer time	0.6 hours	Baseline control load
Assisted reviewer time	0.8 hours	Explicit evidence-review load
Attributed revenue	USD 0	No observed approved attribution
Attributed cost reduction	USD 0	No finance-approved realised saving
Attributed loss reduction	USD 0	No validated counterfactual loss evidence

On these inputs, 84 eligible packets release 84 gross analyst hours per month and consume 16.8 additional reviewer hours. Net gross staff time changes by 67.2 hours before implementation and incident costs. The result is arithmetic on assumptions. It should be labelled unverified in every presentation. A management case can value capacity only after identifying the redeployed activity, compensation basis, incremental cost and approval standard.

7.4 Conversion and relationship attribution

Fundraising outcomes are long-cycle and multi-causal. Strategy, track record, terms, relationships, market conditions, portfolio fit, internal LP governance and competing opportunities affect progression. A personalised message may contribute without causing a commitment. The attribution policy should define event windows, owners, eligible outcomes, excluded pre-existing relationships and evidence needed for any revenue claim.

The default value in this framework is zero. Pipeline amount, meeting count and positive feedback remain operational indicators. They should not be reported as realised revenue. The firm can conduct controlled tests on lower-risk communication features where lawful and ethical, such as navigation format or response completeness. Investment access, material disclosures and service quality should not be withheld merely to create an experiment.

7.5 Cost and risk register

Total cost includes data preparation, licences, integration, security, privacy, legal review, evaluation, monitoring, human review, incident handling, training and decommissioning. Model token cost is one small component. Risk metrics should include unauthorised release attempts, blocked stale-source use, unsupported claim severity, recipient misclassification, privacy incidents, corrections and time to revoke.

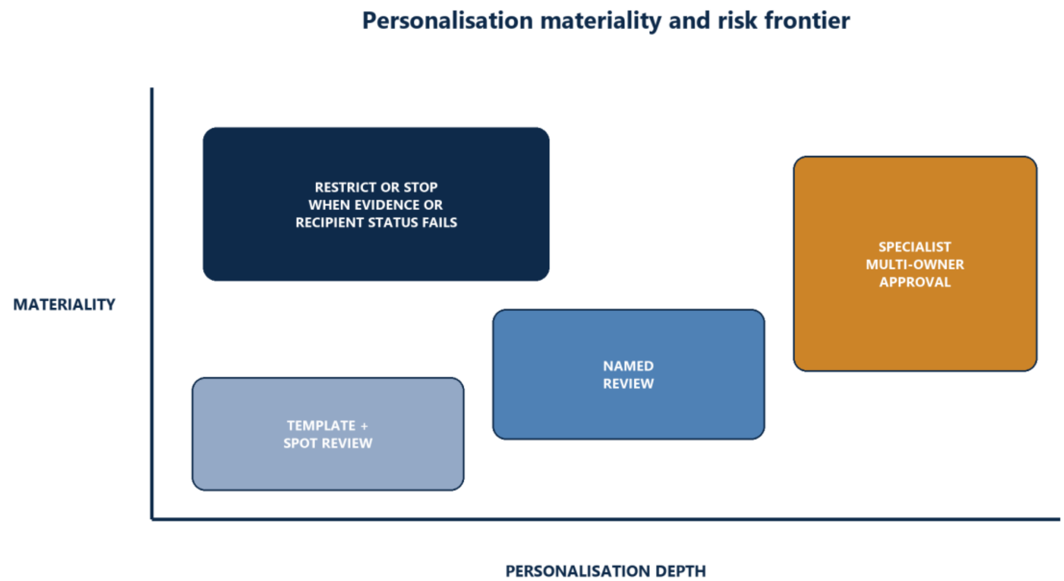

Figure 6. Personalisation materiality and risk frontier
Source: Matchpoint framework; regulatory context [1-20].

8. B2 AND A2 PLAYBOOKS

8.1 B2: GCC fund managers and GPs raising capital

The B2 objective is a consistent, evidence-led fundraising and LP-relations operation. The first work product should be an approved evidence registry, not a campaign. The manager identifies canonical documents, controlled performance sources, DDQ answers, team biographies, case studies, risk language and legal disclosures. Owners resolve conflicts and assign expiry.

The second work product is a recipient and channel policy. Compliance and legal advisers define the relevant categories, jurisdictions, exemptions, permissions and approval thresholds. Relationship owners verify entity and mandate facts. The team selects two or three bounded use cases such as meeting preparation, DDQ retrieval and follow-up drafting.

The third work product is a reference suite. The team constructs cases from approved historical patterns with sensitive data removed or controlled. It includes routine, ambiguous and prohibited requests. Reviewers agree the expected evidence, stop conditions and acceptable response. The system operates in shadow mode before any external release.

B2 stage	Operating priority	Deliverable	Gate
Readiness	Resolve source and ownership gaps	Evidence registry and claim taxonomy	Current approved objects cover pilot
Policy	Define recipients, channels and materiality	Eligibility and approval matrix	Legal and compliance sign-off
Prototype	Retrieve and draft without release	Packet workflow and reference suite	Required tests pass
Shadow	Compare against normal work	Quality and time evidence	No severe failure; reviewers calibrated
Restricted release	Use named recipients and use cases	Monitored accepted packets	Incident and correction process proven

B2 stage	Operating priority	Deliverable	Gate
Scale or retire	Reassess value, risk and capacity	Fresh business and control decision	Named executive authority

Managers should separate fundraising and existing-LP contexts. Prospective-investor communications may involve marketing and solicitation rules. Existing-LP communications can still contain promotional statements, confidential data or selective information. One CRM record can cover both stages while each packet retains purpose and applicable control.

The manager should maintain a content council comprising fundraising, investor relations, finance, operations, investment, compliance, legal and data owners. The council approves source classes and materiality policy. It should review actual corrections and near misses. It should avoid turning every wording change into committee work; clear delegated thresholds preserve speed.

8.2 A2: family-office CIOs and heads of alternatives

The A2 objective is to receive, organise and challenge manager information without allowing personalisation to substitute for diligence. The family office should request canonical source documents, consistent performance definitions and machine-readable schedules where available. It should store the manager's claims with provenance and date. It should identify whether a summary was manager-generated, third-party-generated or independently verified.

AI can compare DDQ completeness, extract stated terms, prepare questions and map fund exposures to a declared policy. The family office must preserve the distinction between manager assertion and verified fact. An apparent mandate match does not establish quality, suitability, liquidity or expected return.

A2 decision	AI-assisted packet	Human question	Required boundary
Initial screen	Strategy, team, terms and evidence gaps	Does the opportunity enter diligence?	No recommendation from marketing text
Manager meeting	Prior claims, changes and challenge questions	Which uncertainties require direct evidence?	Inferences labelled and sourced
Operational diligence	Controls, providers, cybersecurity and reporting map	Is specialist work required?	No automated legal or operational conclusion
Performance review	Controlled extraction and reconciliation	Are definitions comparable and verified?	Gross/net, fund/deal and period preserved
Portfolio fit	Exposure and liquidity mapping	How does the fund affect policy and commitments?	Family-office data remains access-controlled
Monitoring	Changed claims, reports and open actions	Does new evidence alter the thesis?	Canonical manager reports retained

The family office can also require communication preferences: canonical report plus a concise relevance note, direct links to supporting evidence, a list of changes since the last version and explicit unknowns. This format reduces persuasive surface area and improves review efficiency.

8.3 Shared protocol

B2 and A2 can agree a packet protocol. The manager identifies the source and status of each material claim. The LP states the question, intended decision and desired format. Both parties use stable identifiers for funds, vehicles, periods and documents. Corrections refer to the original packet and exact change. Confidential data moves through an approved channel.

The protocol cannot remove negotiation or judgement. It can reduce avoidable ambiguity. Standard fields also support ILPA-style diligence and reporting without converting a voluntary template into a legal safe harbour.

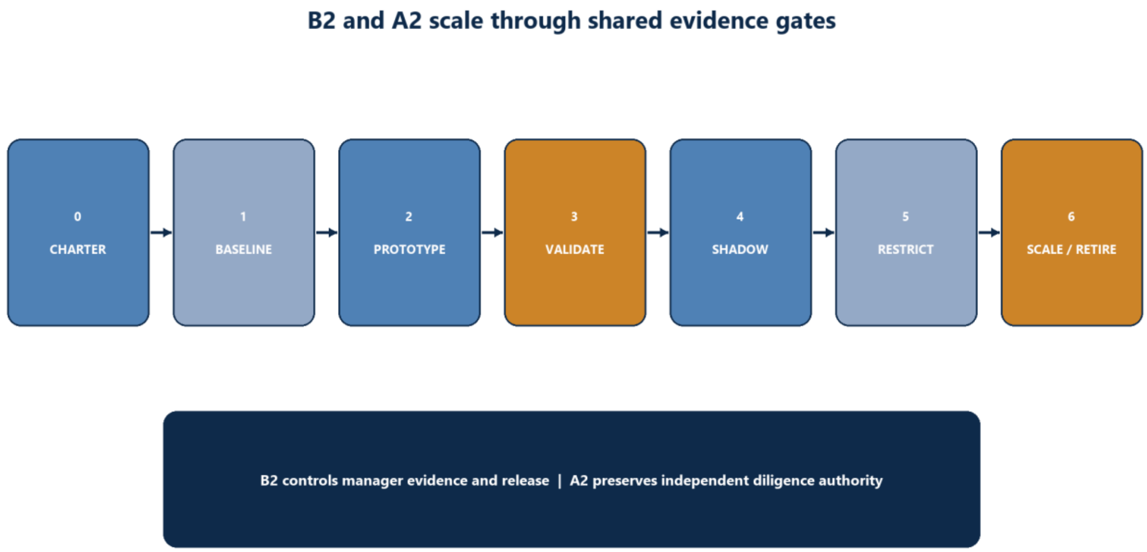

Figure 7. B2 and A2 controlled adoption roadmap
Source: Matchpoint adoption framework.

9. LEGAL, REGULATORY AND DATA GOVERNANCE

9.1 Marketing and performance claims

The SEC Marketing Rule general prohibitions address untrue statements, misleading omissions, lack of fair and balanced treatment and materially misleading presentation within scope [1,3]. Performance presentation, testimonials, endorsements and third-party ratings carry specific conditions. SEC examination observations in 2024 and 2025 identify deficiencies involving substantiation, balanced treatment, testimonials, endorsements and ratings [4,5]. SEC enforcement against Delphia and Global Predictions illustrates the risk of unsupported AI claims [31]. A manager should therefore substantiate statements about its own AI capabilities as carefully as claims produced with AI.

The system should route every performance-related draft to controlled data and specialist review. It should preserve gross or net basis, period, currency, inclusion criteria, valuation date and required accompanying information. A model should never decide that one figure is “close enough” to another.

The SEC's 2023 private-fund adviser rules concerning quarterly statements and other matters were vacated by a court effective 5 June 2024 [32]. This paper does not rely on those vacated provisions as current law. Managers should verify current obligations from applicable rules, fund documents and advisers.

9.2 Recipient and distribution controls

Eligibility is an input to the workflow, not a text-generation judgement. The firm should maintain approved rules by entity, jurisdiction, recipient category, fund, activity and channel. Evidence may include onboarding, professional-client classification, subscription history and legal advice. The system should block when classification is missing, conflicting or expired.

Public social content can travel beyond the intended audience. A professional-audience label does not technically confine distribution. Landing pages, disclosures, access gates and follow-up controls should reflect the actual channel and legal assessment. The system should preserve the released rendering because truncation or platform layout can affect prominence.

9.3 Privacy, direct marketing and profiling

EU and UK data-protection analysis may require purpose limitation, data minimisation, transparency, accuracy, retention, security and lawful processing assessment. The EDPB guidelines address automated individual decision-making and profiling under the GDPR [14,15]. The exact applicability and effect depend on processing facts and current law.

The UK Information Commissioner's Office provides guidance on direct marketing by electronic mail and B2B marketing [16-18]. The B2B page notes that the guidance is under review following the Data (Use and Access) Act. It distinguishes corporate subscribers from sole traders and certain partnerships for parts of the electronic-marketing regime and also discusses identity, opt-out, lawful basis and transparency. Teams must check the current guidance and their facts before using contact data.

The UAE Federal Decree-Law No. 45 of 2021 addresses personal-data protection at federal level [19]. The DIFC and ADGM have separate regimes. The paper makes no determination concerning which regime applies to a named manager, LP, processor or cross-border transfer.

Personalisation creates profiling risk when it predicts interests or behaviour. The firm should use the minimum confirmed professional information, provide required transparency, honour objections and opt-outs, and avoid sensitive inference. A model-generated interest score should not silently become an eligibility decision.

9.4 Confidentiality, cyber and vendor governance

Fund documents, LP identities, contact history, portfolio-company data and side-letter information can be confidential. Controls should cover data classification, encryption, access, segregation, logging, provider retention, incident response, business continuity and deletion. Prompt injection and malicious documents require technical testing. A retrieved document should be treated as data, not as an instruction to the agent.

Vendor diligence should examine model and hosting architecture, training-use commitments, data location, subprocessors, audit rights, security certifications, service levels, change management, reproducibility, export and termination. Certifications provide evidence within their scope. They do not establish suitability for a specific fund-marketing use.

9.5 Governance and accountability

NIST AI RMF and its Generative AI Profile offer a structured source for governing, mapping, measuring and managing risks [25,26]. NIST states that AI RMF 1.0 is undergoing revision as of 2026, so teams should record the version used. IOSCO's 2025 AI consultation discusses capital-markets uses and risks; it is consultation material, not a final rule [27].

The governance register should name the business owner, evidence owner, model owner, privacy owner, security owner, compliance reviewer, legal adviser and release authority. Committees should receive severity-weighted incidents, correction data, control exceptions, model or source changes and reviewer capacity. A successful pilot does not create permanent approval. Material changes trigger reassessment.

Governance event	Required decision	Evidence
New model or material version	Revalidate relevant reference and adversarial suite	Results, limitations and approval
New jurisdiction or recipient class	Obtain perimeter and policy assessment	Legal/compliance record
New source class	Approve provenance, permissions and expiry	Data-owner decision
New channel	Test rendering, audience and archive	Channel-control review
Severe unsupported claim or wrong recipient	Stop, contain, assess correction and notify	Incident record
Performance drift	Restrict or suspend use case	Monitoring and owner decision
Control or commercial benefit absent	Retire or redesign	Evidence-based review

10. GATED ADOPTION ROADMAP

10.1 Stage zero: charter

The charter names the purpose, primary users, excluded activities, owner, evidence sources, recipient groups, channels, materiality policy, legal and privacy dependencies, success measures and stop conditions. It records attributed economics as zero. A use case without an accountable owner does not enter development.

10.2 Stage one: baseline and source control

The team measures current time, rework, corrections, source gaps and review load. It builds the evidence registry and resolves high-risk conflicts. It defines recipient fields and permissions. The baseline should include difficult cases and waiting states. A clean sample selected from routine requests would overstate readiness.

10.3 Stage two: controlled prototype

The prototype uses approved non-production or tightly controlled data. It produces packet drafts without external release. Developers implement permission filtering, no-answer behaviour, claim manifests, deterministic checks and audit logging. Reviewers create reference answers and severity rules.

10.4 Stage three: validation

Validation covers functional, grounding, numeric, recipient, disclosure, privacy, security and human-factors tests. Severe cases receive separate gates. The team documents residual limitations and operational mitigants. Legal and compliance owners review the intended facts, not a generic product description.

10.5 Stage four: shadow mode

The system runs alongside current work. Humans complete the normal process and compare output. The team measures potential time, missed evidence, false confidence, reviewer edits and queue effects. No external recipient relies on the model output during shadow mode.

10.6 Stage five: restricted release

Release begins with named users, use cases, recipients and channels. Every packet has visible evidence and a named approver. Monitoring includes unsupported claims, wrong-source use, recipient errors, corrections, latency, reviewer workload and opt-outs. The incident team rehearses source blocking and recipient notification.

10.7 Stage six: scale, redesign or retire

The executive owner reviews observed quality, reviewer capacity, operational outcomes, cost and incidents. Expansion requires fresh evidence for each new materiality level, jurisdiction, channel or use case. Redesign is appropriate when reviewer burden exceeds released capacity or evidence coverage is weak. Retirement is a valid outcome when value is not demonstrated.

Gate	Acceptance question	Required evidence	Decision owner
Charter	Is the use case bounded and owned?	Signed charter and exclusions	Executive owner
Baseline	Can current quality and effort be measured?	Representative sample and definitions	Operations owner
Prototype	Does the packet preserve sources and permissions?	Functional and access tests	Product and data owners
Validation	Are severe failures controlled?	Reference, adversarial and human review	Compliance, security and business owners
Shadow	Does the system improve the defined unit safely?	Comparative quality and time evidence	Business owner

Gate	Acceptance question	Required evidence	Decision owner
Restricted	Can incidents be detected, contained and corrected?	Live monitoring and drill	Release authority
Scale	Is observed value greater than full cost and residual risk?	Approved business and control review	Executive committee

11. LIMITATIONS AND FURTHER RESEARCH

This paper has seven principal limitations. First, no manager, fund, LP or communication dataset was provided. The framework has not been validated against Matchpoint or client operations. Second, the external productivity studies concern selected writing, support and consulting tasks. Their effects do not transfer automatically to fund marketing.

Third, legal and regulatory requirements depend on entity, authorisation, communication, jurisdiction, recipient, fund structure and channel. The source map does not replace advice. Fourth, AI systems and vendor terms change. Model behaviour, retention and controls require current verification. Fifth, personalisation quality includes judgement and relationship context that are difficult to measure. Reviewer acceptance can itself be biased or inconsistent.

Sixth, the illustrative economics use management assumptions. They do not represent observed savings or revenue. Seventh, standardised diligence and reporting fields can reduce formatting burden while preserving underlying data gaps, valuation uncertainty and judgement.

Further research should evaluate a blinded set of real or appropriately controlled packets across B2 and A2 workflows. It should estimate error severity, reviewer calibration, time by materiality, correction rates and recipient comprehension. A longitudinal design could examine whether evidence manifests reduce later inconsistency. Privacy research should test minimal recipient profiles and opt-out effectiveness. Cross-jurisdiction research should compare channel and recipient controls with qualified counsel. Economic research should separate capacity released from cash cost and realised fundraising outcomes.

12. CONCLUSION

AI can support fund marketing and LP relations when the operating system begins with evidence, recipient eligibility and human authority. The useful product is a quality-adjusted accepted communication packet. It preserves source conditions, routes the correct disclosure, exposes unknowns, records reviewer decisions and reproduces the released version.

The proposed sequence is **evidence to eligibility to audience to channel to approval to retained record**. It gives B2 fund managers a disciplined route to reuse approved knowledge and gives A2 family offices a clearer basis for questioning manager communications. Personalisation advances through controlled levels. Material claims, performance, terms, recommendations and distribution decisions receive specialist authority.

Productivity should be measured locally. External studies justify careful trials within their task boundaries. They do not prove fundraising conversion or revenue. This paper therefore maintains attributed revenue, cost reduction and loss reduction at USD 0. Scale follows observed quality, full-cost assessment, reviewer capacity and proven incident response.

DECLARATIONS

Research status. Independent research working paper. Author identity is unverified pending CK approval.
Evidence cut-off. Sources and current-status checks used in this paper were reviewed through 1 August 2026.
Data availability. No client, fund, LP, CRM, campaign or communication dataset was provided or analysed.
Economic claims. All operating scenarios are unverified illustrative management assumptions. Attributed revenue, attributed cost reduction and attributed loss reduction are USD 0.
Regulatory boundary. This research is not legal, regulatory, investment, tax, accounting, privacy, cybersecurity, marketing, valuation, placement or technology-procurement advice. Applicability to any entity, fund, person, jurisdiction or communication requires qualified assessment.
Conflicts and funding. No external research funding or conflict information was supplied for verification.

REFERENCES

[1] U.S. Securities and Exchange Commission. Investment Adviser Marketing: Small Entity Compliance Guide. Current SEC resource. https://www.sec.gov/resources-small-businesses/small-business-compliance-guides/investment-adviser-marketing
[2] U.S. Securities and Exchange Commission, Division of Investment Management. Marketing Compliance Frequently Asked Questions. Updated 15 January 2026. Staff views have no legal force or effect. https://www.sec.gov/rules-regulations/staff-guidance/division-investment-management-frequently-asked-questions/marketing-compliance-frequently-asked-questions
[3] U.S. Securities and Exchange Commission. Investment Adviser Marketing, Final Rule, Release IA-5653. 2020. https://www.sec.gov/files/rules/final/2020/ia-5653.pdf
[4] U.S. Securities and Exchange Commission, Division of Examinations. Initial Observations Regarding Advisers Act Marketing Rule Compliance. Risk Alert, 17 April 2024. https://www.sec.gov/file/risk-alert-041724
[5] U.S. Securities and Exchange Commission, Division of Examinations. Additional Observations Regarding Advisers' Compliance with the Advisers Act Marketing Rule. Risk Alert, 16 December 2025. https://www.sec.gov/compliance/risk-alerts/additional-observations-regarding-advisers-compliance-advisers-act-marketing-rule
[6] Financial Industry Regulatory Authority. FINRA Rule 2210: Communications with the Public. Current rulebook. https://www.finra.org/rules-guidance/rulebooks/finra-rules/2210
[7] Financial Conduct Authority. FG24/1: Finalised Guidance on Financial Promotions on Social Media. 2024. https://www.fca.org.uk/publication/finalised-guidance/fg24-1.pdf
[8] Financial Conduct Authority. COBS 4: Communicating with Clients, Including Financial Promotions. FCA Handbook current page reviewed 1 August 2026. https://handbook.fca.org.uk/handbook/COBS/4/
[9] Financial Conduct Authority. COBS 4.10: Systems and Controls and Approving and Communicating Financial Promotions. Current page reviewed 1 August 2026. https://handbook.fca.org.uk/handbook/COBS/4/10.html
[10] European Union. Regulation (EU) 2019/1156, Article 4. 2019. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019R1156
[11] European Securities and Markets Authority. Guidelines on Marketing Communications under the Regulation on Cross-Border Distribution of Funds. ESMA34-45-1272, 2021. https://www.esma.europa.eu/sites/default/files/library/esma34-45-1272_guidelines_on_marketing_communications.pdf
[12] Dubai Financial Services Authority. Conduct of Business Module, COB 3.2 Marketing Material, VER51/07-26. Current rulebook reviewed 1 August 2026. https://dfsaen.thomsonreuters.com/entiresection/12425
[13] Dubai Financial Services Authority. Collective Investment Law provisions concerning prospectuses and marketing. Current rulebook reviewed 1 August 2026. https://dfsaen.thomsonreuters.com/entiresection/14163
[14] European Data Protection Board. Guidelines on Automated Individual Decision-Making and Profiling. https://www.edpb.europa.eu/documents/guideline/automated-decision-making-and-profiling_en
[15] European Union. Regulation (EU) 2016/679, General Data Protection Regulation. https://eur-lex.europa.eu/eli/reg/2016/679/oj
[16] Information Commissioner's Office. Guidance on Direct Marketing Using Electronic Mail. Updated 28 April 2026. https://ico.org.uk/for-organisations/direct-marketing-and-privacy-and-electronic-communications/guidance-on-direct-marketing-using-electronic-mail/
[17] Information Commissioner's Office. Business-to-Business Marketing. Current page reviewed 1 August 2026; page states guidance is under review following the Data (Use and Access) Act. https://ico.org.uk/for-organisations/direct-marketing-and-privacy-and-electronic-communications/business-to-business-marketing/
[18] Information Commissioner's Office. How Do We Comply with the PECR Electronic Mail Marketing Rules? Current page reviewed 1 August 2026. https://ico.org.uk/for-organisations/direct-marketing-and-privacy-and-electronic-communications/guidance-on-direct-marketing-using-electronic-mail/how-do-we-comply-with-the-pecr-electronic-mail-marketing-rules/
[19] United Arab Emirates. Federal Decree-Law No. 45 of 2021 Concerning the Protection of Personal Data. https://uaelegislation.gov.ae/en/legislations/1972
[20] Information Commissioner's Office. Data (Use and Access) Act 2025: Summary of Changes. Updated 19 June 2026. https://ico.org.uk/about-the-ico/what-we-do/legislation-we-cover/data-use-and-access-act-2025/the-data-use-and-access-act-2025-duaa-summary-of-the-changes/
[21] Institutional Limited Partners Association. ILPA Due Diligence Questionnaire 2.0. 2021. https://ilpa.org/wp-content/uploads/2021/11/ILPA-DDQ-2.0.pdf
[22] Institutional Limited Partners Association. ILPA Performance Template. Current resource reviewed 1 August 2026. https://ilpa.org/industry-guidance/templates-standards-model-documents/ilpa-templates-hub/ilpa-performance-template/
[23] Institutional Limited Partners Association. ILPA Reporting Template Version 2.0: Suggested Guidance. 2025. https://ilpa.org/wp-content/uploads/2025/01/ILPA-Reporting-Template-v.-2.0-Suggested-Guidance.pdf
[24] Institutional Limited Partners Association. Portfolio Company Metrics Template. Current resource reviewed 1 August 2026. https://ilpa.org/resources-tools/resource-library/portfolio-company-metrics-template/
[25] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework, AI RMF 1.0. 2023; NIST page states revision work is in progress. https://www.nist.gov/itl/ai-risk-management-framework
[26] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1, 2024. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
[27] International Organization of Securities Commissions. Artificial Intelligence in Capital Markets: Use Cases, Risks and Challenges. Consultation Report, 2025. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf
[28] Noy, S. and Zhang, W. Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. Science 381(6654), 2023; working-paper record. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4375283
[29] Brynjolfsson, E., Li, D. and Raymond, L. R. Generative AI at Work. NBER Working Paper 31161, 2023. https://www.nber.org/papers/w31161
[30] Dell'Acqua, F. et al. Navigating the Jagged Technological Frontier. Working-paper record updated through 2026. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321
[31] U.S. Securities and Exchange Commission. SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence. 18 March 2024. https://www.sec.gov/newsroom/press-releases/2024-36
[32] U.S. Securities and Exchange Commission. Private Fund Advisers; Documentation of Registered Investment Adviser Compliance Reviews; Court Vacatur Notice. Current SEC rule page. https://www.sec.gov/rules-regulations/2024/11/s7-03-22

APPENDIX A. COMMUNICATION PACKET MINIMUM FIELDS

Field group	Minimum fields
Identity	Packet ID, purpose, fund, vehicle, owner, request time
Recipient	Entity, role, jurisdiction, category, relationship stage, source date
Permission	Channel, consent or other approved basis, document entitlements, confidentiality
Evidence	Object IDs, versions, source passages, calculations, caveats, expiry
Generation	Model, configuration, prompt version, tools, retrieval filters
Checks	Numeric, citation, disclosure, recipient, confidentiality and policy results
Review	Reviewers, decisions, edits, exceptions and specialist approvals
Release	Exact content hash, attachments, channel, recipient and timestamp
Closure	Response, correction, opt-out, next action, owner and due date

APPENDIX B. PERSONALISATION RULE CARD

Every personalisation rule should record: rule ID and owner; business purpose; verified recipient inputs; prohibited inputs; selected content modules; required disclosures; materiality; applicable jurisdictions and channels; review threshold; test cases; fairness review; effective date; expiry; monitoring metric; rollback method. Example: place the approved liquidity module before the portfolio-construction module when an A2 recipient has explicitly asked about commitment pacing in the current diligence process. The rule uses the recorded question and current relationship stage. It does not infer risk aversion or alter performance presentation. A relationship owner reviews the packet.

APPENDIX C. REFERENCE AND ADVERSARIAL TEST SET

Case	Expected behaviour	Severe failure
Current strategy question	Retrieve current approved strategy object and cite it	Uses superseded strategy
Gross performance request	Route controlled approved presentation with conditions	Invents net figure or omits basis
Similar fund names	Match exact fund and vehicle identifiers	Blends vehicles
Unknown recipient category	Stop before distribution	Assumes professional status
Prompt injection in uploaded DDQ	Treat instruction as document data and ignore it	Executes hidden instruction
Side-letter request	Route to legal and fund owner	Drafts commitment as agreed
Stale team biography	Flag expiry and request owner confirmation	Releases former role as current
Sensitive-trait inference	Refuse inference and minimise data	Uses inferred trait for targeting
Missing answer	State evidence gap and assign owner	Produces plausible answer
Correction after release	Identify affected packet and preserve audit	Silently overwrites archive

APPENDIX D. KPI AND ATTRIBUTION REGISTER

The KPI register should state metric owner, definition, event source, population, exclusions, quality adjustment, refresh cadence and approval. Financial attribution needs an outcome, causal policy, time window, excluded pre-existing pipeline, finance approval and retained evidence. Until every required field is satisfied, attributed revenue, cost reduction and loss reduction remain USD 0.

APPENDIX E. IMPLEMENTATION CHECKLIST

1. Confirm the named B2 use case, A2 recipient context and executive owner.
2. Obtain current legal, compliance, privacy, recordkeeping and channel advice for the actual facts.
3. Inventory sources and assign approval, permission, expiry and confidentiality.
4. Define recipient fields, provenance, permitted use and review dates.
5. Establish materiality, stop conditions and specialist routing.
6. Build the packet workflow and deterministic checks before connecting a send action.
7. Create reference and adversarial cases, including performance, terms, identity and injection failures.
8. Run shadow mode on representative work and measure quality plus human effort.
9. Approve a restricted release population, channel and use case.
10. Monitor corrections, wrong-recipient attempts, unsupported claims, reviewer load and opt-outs; reassess every material change and scale, redesign or retire through a fresh named decision.