MatchPoint Research Working Paper No. 2026/T34

The AI-Native Family Office: Deal Screening for Two-Person Teams

A controlled operating model for intake, diligence extraction, committee-memo drafting and accountable investment decisions

Prepared by Matchpoint Partners
Author attribution pending CK approval
August 2026

ABSTRACT

Background. A compact family-office team must screen a diverse opportunity set while preserving mandate, permission, diligence and decision evidence. **Objective.** This paper develops an AI-native deal-screening operating model for A7 GCC family offices and ultra-high-net-worth asset owners and A2 family-office CIOs and heads of alternatives. **Approach.** The analysis reviews 43 primary and authoritative sources available through 2 August 2026 and joins immutable sources, opportunity and evidence objects, deterministic mandate gates, permissioned retrieval, Claude-assisted extraction and drafting, and named human approvals. **Findings.** Claude can support bounded classification, extraction, comparison, retrieval and drafting when source custody, calculations, permissions, specialist determinations and investment decisions remain under approved controls. The two-person-versus-ten-person statement is an unverified operating hypothesis; no supplied family-office study proves the equivalence. **Implications.** Throughput should be measured as quality-gated decision-ready packs using comparable cases, observed effort, failure-aware quality and approved cost evidence. Worked cases and operating inputs are unverified illustrative scenarios; attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence exists.

Highlights

Primary users. A7 GCC family offices and A2 family-office CIOs and heads of alternatives.

Evidence unit. An atomic proposition with source, period, perimeter, permission, transformation, reviewer and refresh trigger.

Architecture. Controlled intake feeds mandate gates, an evidence graph, bounded Claude tasks, deterministic tools, exceptions and authorised workflow.

Decision boundary. The model supports evidence work; named professionals, specialists, principals and committees retain consequential decisions.

Economics. Attributed revenue, cash saving, loss reduction and alpha remain USD 0 until observed evidence is approved.

JEL Classification: D81, D91, G11, G23, G24, G32, G34, K22, M15, O33

Keywords: GCC family office, deal screening, investment diligence, Claude, investment committee memo, evidence graph, beneficial ownership, mandate fit, suitability, agentic workflows

1. INTRODUCTION

A compact family-office investment team faces an operating problem with several dimensions. It must source widely enough to preserve opportunity, screen quickly enough to protect senior attention, investigate deeply enough to make an accountable decision, and retain a record that survives staff turnover, family-governance review and later portfolio monitoring. A two-person team can receive teasers, decks, models, data-room notices, adviser emails, fund questionnaires and direct-investment proposals across many sectors and structures. Volume alone says little about whether the team can reach a reviewable investment decision.

This paper develops an AI-native deal-screening operating model for A7 GCC family offices and ultra-high-net-worth asset owners, and A2 family-office chief investment officers and heads of alternatives. Claude-based extraction, comparison and drafting are treated as bounded components inside a controlled evidence system. The system preserves original files, normalises opportunity data, applies approved mandate and suitability gates, assembles diligence evidence, routes exceptions, and drafts a cited investment-committee memo. Principals, investment professionals and qualified specialists retain all investment, legal, tax, compliance, valuation and governance decisions.

The Topic Tracker states a provocative proposition: a two-person investment team with Claude-based screening, diligence extraction and investment-committee-memo drafting can match the throughput of a ten-person institutional team. **[Unverified]** No supplied family-office study, current primary source or observed Matchpoint deployment establishes that equivalence. This paper converts the proposition into a testable operating hypothesis. Throughput is defined at the point where a complete, cited and reviewable decision pack reaches the authorised decision maker. A count of summaries, extracted fields or drafted pages is not decision throughput.

Research question	Operating answer
What can an AI-native workflow support?	Permissioned intake, classification, extraction, comparison, retrieval, exception assembly, workflow routing and cited drafting within approved schemas.
What must remain controlled?	Source custody, arithmetic, mandate rules, access rights, versioning, approval gates, records, suitability determinations and consequential decisions.
What counts as throughput?	Opportunities that reach a defined decision state with required evidence, quality gates, named review and a reproducible record.
How is diligence evidence represented?	As atomic evidence objects linked to opportunities, entities, claims, documents, analyses, exceptions, decisions and later monitoring.
How is an IC memo generated?	From approved evidence objects and analysis records, with sentence-level support, disclosed gaps and retained human authorship of the recommendation.
What proves a headcount or cost claim?	A preregistered, like-for-like pilot using comparable cases, quality thresholds, loaded cost inputs, observed hours and approved outcomes.

The legal and governance context matters. UAE federal family-business legislation provides a framework for registered family businesses and family governance [1]. Dubai legislation established a Family Companies Centre and regulates family property arrangements within its scope [2-3]. DIFC and ADGM provide distinct regimes and structures relevant to family arrangements and family offices [4-8]. Data protection, beneficial ownership, anti-money-laundering, electronic-transactions and evidence requirements affect document handling and entity diligence [9-21]. Those sources do not create a universal family-office operating model. Applicability depends on the entity, structure, activity, jurisdiction and facts, and qualified advisers determine legal effect.

The paper also uses authoritative investment-practice and AI-governance materials. ILPA diligence resources can structure private-fund review [23-25]. CFA Institute standards support diligence, suitability and loyalty principles for investment professionals subject to those standards [26-28]. Anthropic materials describe agent patterns, citations, evaluations, tool use and autonomy [29-36]. NIST and W3C materials support AI-risk and provenance design [37-40]. Cross-domain productivity studies provide external benchmarks [41-43]. They do not prove family-office investment-team productivity.

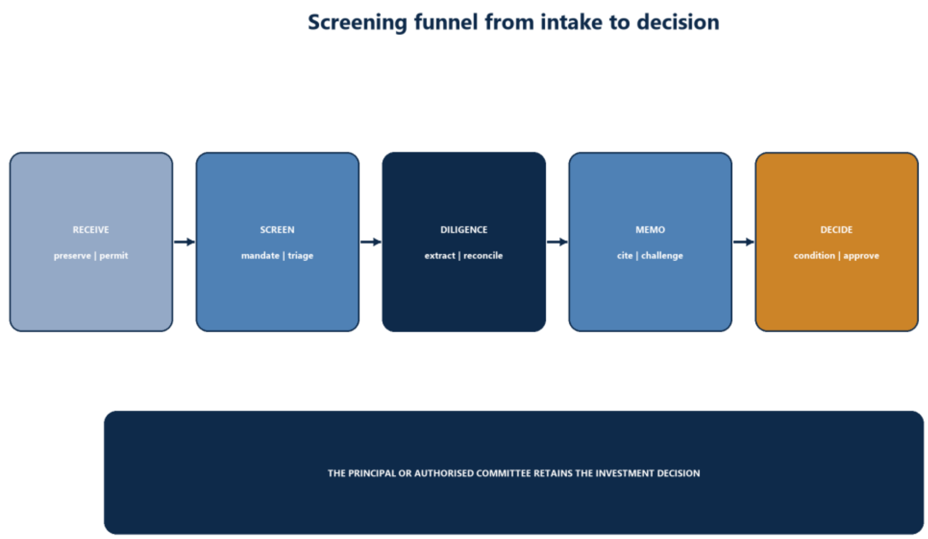

Figure 1. Screening funnel from permissioned intake to accountable decision

Source: Matchpoint family-office screening framework; governance context [1-8,23-28].

The paper provides a reference architecture, operating controls, schemas, tables, workflows and one **[Unverified illustrative scenario]**. It does not provide investment, legal, regulatory, tax, accounting, valuation, privacy, cybersecurity or technology advice. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence is supplied.

2. A7 AND A2 DECISION PERIMETER

2.1 A7 GCC family offices and ultra-high-net-worth asset owners

A7 represents principals, owners and family-office teams allocating family capital across public and private markets, operating businesses, direct investments, co-investments, funds, real assets and strategic opportunities. The mandate may include financial return, capital preservation, liquidity, family control, reputation, succession, strategic access, geographic diversification and impact. Those objectives can conflict. A screening system must preserve the family's authorised hierarchy of objectives instead of replacing it with a generic score.

The family-office form also varies. A single-family office may serve one family through one or more legal entities. A multi-family office serves multiple families and may be subject to different regulatory and contractual requirements. The SEC family-office rule and its staff FAQ illustrate a specific United States exclusion with defined conditions [22]. DIFC and ADGM maintain their own current regimes and guidance [4-8,20-21]. A label such as family office does not determine regulatory status. The system must record the actual entity, activity, jurisdiction and authorisation position supplied by qualified advisers.

2.2 A2 family-office CIOs and heads of alternatives

A2 represents investment professionals responsible for translating family objectives into portfolio and transaction decisions. Their operating duties may include pipeline management, manager selection, direct-investment underwriting, portfolio construction, cash planning, co-investment coordination, diligence management, committee preparation and monitoring. The system should serve the CIO as an evidence and workflow layer. It should not create an independent investment mandate or silently change approved policy.

In a compact office, roles can overlap. One professional may source, analyse and draft; another may challenge, coordinate advisers and prepare a recommendation for a principal or committee. Overlap increases the need for visible decision rights. A model can perform repeated transformations. It cannot supply independence, fiduciary judgement, family authority or specialist accountability.

2.3 Decision-rights map

Decision	AI-native contribution	Retained authority
admit an opportunity to the funnel	classify source, create record, identify duplicates and test minimum fields	authorised investment professional
determine mandate fit	apply approved eligibility rules and show evidence or missing fields	CIO, principal or authorised policy owner
accept identity and ownership	extract and reconcile identifiers and ownership claims	compliance, legal and authorised onboarding owners
determine suitability	assemble portfolio, liquidity, risk and restriction evidence	authorised investment decision maker under applicable duties
select diligence workstreams	propose a template from deal type and known risks	deal lead and relevant specialists
accept a valuation	assemble valuation inputs, methods and exceptions	authorised valuation and investment authority
accept legal or tax consequences	retrieve sources and frame questions	qualified legal and tax advisers
recommend an investment	draft a cited synthesis from approved evidence	named investment professional
approve or reject	record decision, conditions and rationale	principal, committee or authorised delegate
communicate externally	prepare a permissioned draft	authorised communications owner

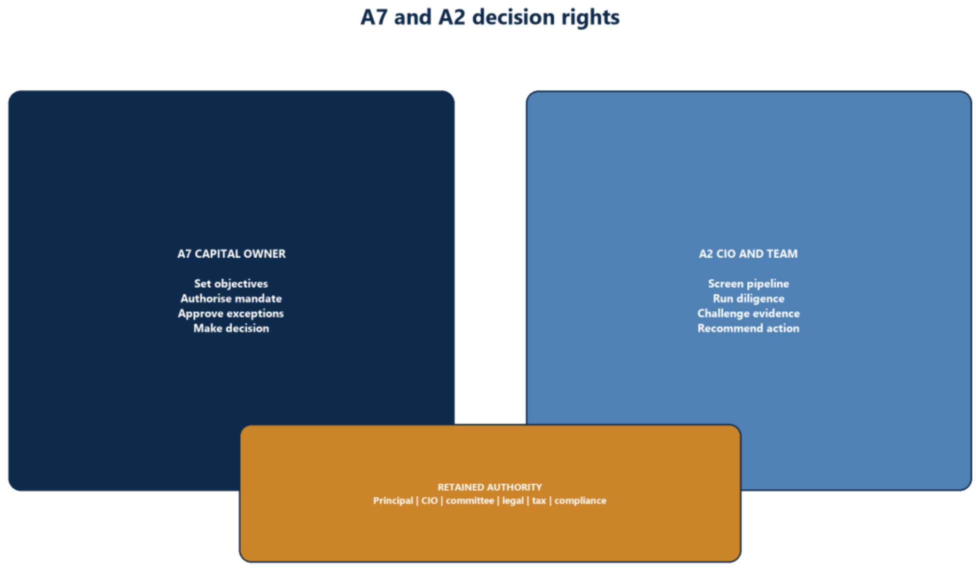

Figure 2. A7 and A2 decision-rights map

Source: Matchpoint ICP and decision-rights framework.

2.4 Operating states

Every opportunity should use a controlled state vocabulary. **Captured** means the original source and sender are retained. **Eligible** means approved minimum mandate gates are supported. **Diligencing** means a named owner, workplan and evidence matrix exist. **Conditioned** means progression depends on named evidence or action. **Decision-ready** means required gates pass and the memo has completed review. **Approved**, **rejected** and **deferred** are authorised decisions. **Archived** preserves a closed record. These states must be distinguishable from model confidence.

3. THE THROUGHPUT HYPOTHESIS AND EVIDENCE BOUNDARY

3.1 Statement under test

The operational hypothesis is: a two-person family-office investment team supported by a controlled AI-native workflow can produce the same number of decision-ready packs per defined period as a specified ten-person comparator while meeting the same quality, risk and governance gates. Each term requires a definition. The family-office team, comparator, opportunity mix, decision pack, period, quality gates and retained specialist work must be stated before measurement.

The hypothesis does not mean that two people possess the combined expertise, independence, relationships or decision rights of ten professionals. It addresses a bounded output under a defined workflow. A result can apply only to the evaluated population and operating conditions.

3.2 External productivity evidence

Three frequently cited studies concern different work domains. An NBER field study reported an average productivity increase for customer-support agents using a generative-AI assistant [41]. A controlled Microsoft and GitHub experiment measured faster completion of a defined software-development task for participants with access to GitHub Copilot [43]. A Harvard Business School study of consultants reported faster and higher-quality performance on tasks inside the model's capability frontier, alongside lower correctness on a task outside that frontier [42]. These are external-domain benchmarks. They do not establish family-office investment productivity, diligence quality, loss prevention or two-person equivalence.

Anthropic's Economic Index, Claude Code expertise research and autonomy research use product telemetry, experiments or model-behaviour analysis within their described scopes [34-36]. They can inform evaluation design. They do not supply a controlled study of family-office deal screening.

3.3 Measurement unit

The primary unit should be a decision-ready opportunity pack, defined by:

1. a verified opportunity and entity record;
2. an approved mandate-fit determination;
3. a complete applicable evidence matrix or disclosed exceptions;
4. reproducible analysis with source and transformation lineage;
5. a cited memo that separates fact, representation, analysis and judgement;
6. review by named authorities;
7. a recorded decision state; and
8. monitoring objects for every approved condition or investment.

Secondary units may include screened opportunities, evidence objects, resolved exceptions, reviewer hours, calendar time and costs. They should never replace the primary quality-gated unit.

3.4 Failure-aware comparison

The measurement design must count rework and material failures. A fast draft that omits a beneficial owner, uses a stale model, confuses currencies, cites an unsupported sentence or breaches permission is a failed output. A comparison should record false eligibility, false rejection, extraction errors, unsupported claims, missed conflicts, arithmetic failures, access incidents, late specialist escalation, reviewer overrides and post-decision corrections.

Measure	Required definition	Evidence
decision-ready throughput	packs passing all gates per period	immutable state transitions and approvals
cycle time	elapsed time from complete intake to decision-ready state	system timestamps with paused states defined
human effort	active minutes by role and task	approved time records or observed study
quality	task-specific correctness and completeness	golden set, blinded review and severity weighting
rework	time and changes after first review	version history and reviewer record
risk outcome	specified process or investment outcome	approved event definition and observed evidence
cost	loaded labour plus approved technology and adviser costs	finance-approved inputs and invoices
comparator parity	same case population, scope and gates	preregistered protocol and case assignment

4. FAMILY-OFFICE GOVERNANCE AND MANDATE MAP

4.1 Governance before automation

UAE Federal Decree-Law No. 37 of 2022 addresses family businesses registered under its scope and permits governance mechanisms contemplated by the law [1]. Dubai Decree No. 45 of 2022 established the Family Companies Centre [2]. Dubai Law No. 9 of 2020 regulates family property within its stated scope [3]. DIFC Family Arrangements Regulations and ADGM family-office and foundation materials provide additional structures and services in their jurisdictions [4-8]. The relevant family, entity and advisers determine which arrangements apply.

An AI-native workflow should receive a signed or otherwise authorised mandate configuration. The configuration is a system input, not a model-generated policy. It should record policy owner, approval date, effective date, version, review date and superseded versions.

4.2 Mandate object

Mandate dimension	Example fields	Control question
capital owner	family entity, trust, foundation, holding company or vehicle	which legal person or arrangement owns the capital?
objective	preservation, return, liquidity, control, strategic access, impact	what is authorised and how are conflicts ranked?
allocation	asset class, geography, currency, vehicle, liquidity and concentration	which limits are hard and which require approval?
transaction	size, stage, instrument, control, co-investment, leverage	which structures are eligible?
counterparty	sponsor, manager, adviser, jurisdiction and restricted party	which parties require enhanced review?
risk	loss tolerance, liquidity horizon, drawdown, complexity and reputation	which risk definitions and evidence apply?
governance	proposer, reviewer, committee, principal and specialist rights	who can recommend, condition, approve or reject?
information	confidentiality, privilege, data location, sharing and retention	who may access each source and output?

4.3 Policy logic

Hard exclusions should be deterministic rules. Examples include prohibited jurisdictions, unauthorised asset classes, minimum liquidity reserves or transaction sizes outside delegated authority. Soft preferences should rank attention without producing a decision. Examples include preferred sectors, relationship value, thematic fit and strategic access. Judgement items should route to a named person. The system should show which policy version was applied and how every rule evaluated.

4.4 Conflicts and loyalty

CFA Institute Standard III(A) addresses loyalty, prudence and care for members and candidates subject to the standard [28]. The family-office setting may involve additional duties under contracts, law, entity governance and professional status. The workflow should collect conflict facts: ownership, board roles, adviser compensation, placement fees, personal relationships, co-investor rights, related-party status and information barriers. A model may flag a possible conflict from supplied evidence. An authorised person determines the conflict treatment.

4.5 Suitability perimeter

CFA Institute Standard III(C) addresses suitability for members and candidates subject to the standard [27]. A family's authorised policy may require comparable analysis even where that standard is not directly applicable. The system should connect each proposed investment to portfolio exposure, liquidity, commitments, concentration, currency, leverage, downside and family constraints. Suitability is a portfolio-level judgement. A deal can be attractive in isolation and unsuitable for the current portfolio.

5. DEAL-FUNNEL OPERATING MODEL

5.1 Funnel stages

The funnel should separate low-cost classification from high-cost judgement. A reference sequence is:

1. receive and preserve;
2. identify source, sender and permission;
3. create or match opportunity and entity records;
4. test minimum eligibility;
5. prioritise for human review;
6. open a diligence workplan;
7. extract and reconcile evidence;
8. complete analysis and specialist review;
9. draft and review the IC memo;
10. record the decision; and
11. promote approved obligations and assumptions into monitoring.

No stage should erase the earlier state. A rejected opportunity remains searchable under its retention policy, with the reasons, evidence and decision authority. Duplicate opportunities should link to one canonical record while preserving each introduction and permission context.

5.2 Service-level definitions

Service levels should distinguish system latency, analyst activity, adviser time and waiting for the counterparty. A response-time target can measure acknowledgement. It cannot prove diligence quality. The system should pause or classify elapsed time when required evidence is unavailable. This makes throughput comparisons fairer and prevents counterparty delay from being mislabelled as team inefficiency.

5.3 Triage without false precision

A triage score may allocate attention, provided every component is visible and the score is never treated as expected return or probability of success. Hard gates remain separate. A sample triage might combine mandate proximity, strategic relevance, information completeness, decision urgency and expected diligence burden. Weights require owner approval and pilot testing. Unverified weights should be labelled and must not enter production.

Triage component	Permitted evidence	Prohibited interpretation
mandate proximity	approved policy fields and supported opportunity fields	investment attractiveness
strategic relevance	authorised themes and relationship facts	expected return
completeness	received versus required fields	counterparty quality
urgency	stated process dates and verified deadlines	pressure to bypass controls
diligence burden	applicable workstreams and document population	reason to lower quality gates
conflict complexity	disclosed relationships and policy rules	resolved conflict

5.4 Exception-first design

Senior time is scarce. The system should present material exceptions before narrative. An exception record states the object, expected condition, observed evidence, discrepancy, materiality basis, affected decision, owner, due date, permissible actions and closure evidence. Grouping exceptions by consequence helps the team distinguish blocking identity questions from low-impact formatting differences.

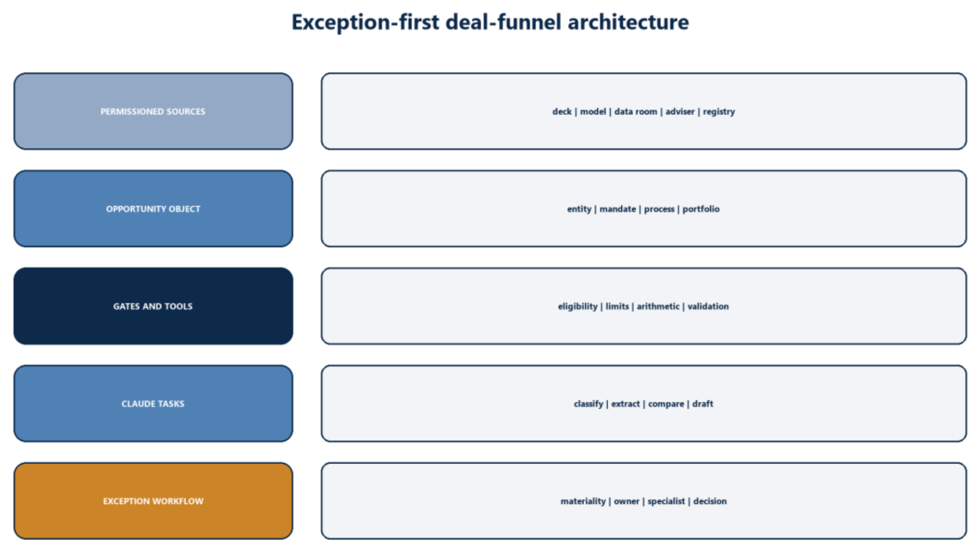

Figure 3. Exception-first deal-funnel architecture

Source: Matchpoint family-office operating architecture; AI workflow context [29-33,37-40].

6. SOURCE INTAKE AND OPPORTUNITY NORMALISATION

6.1 Permissioned intake

Each received item should be captured with sender, recipient, timestamp, channel, original filename, file hash, malware-scan result, declared confidentiality, legal privilege status where supplied, permitted users, permitted purpose, retention class and deletion trigger. The original remains immutable. Derived text, tables and images are separate objects with provenance.

UAE personal-data protection legislation establishes requirements for processing personal data within its scope [14]. DIFC and ADGM have their own data-protection regimes and guidance [9,15]. The correct regime and cross-border position depend on facts and jurisdiction. The workflow should store the documented legal basis or approved processing instruction supplied by the data owner and advisers. It should not infer permission from file availability.

6.2 Opportunity object

The opportunity object links the commercial proposition to its legal and data perimeter. It should include opportunity identifier, source, relationship owner, sponsor or manager, target or vehicle, instrument, stage, amount range, currency, geography, sector, process dates, requested decision, confidentiality class, known advisers and linked source files. Unknown fields remain null.

6.3 Entity resolution

Entity resolution should prefer durable identifiers: registration number, legal-entity identifier where available, official registry record, tax number where permitted, and controlled internal identifier. Names, websites and addresses are supporting attributes. A fuzzy name match is a candidate relationship, never proof.

The graph should separate legal ownership, beneficial ownership, control, management, board representation, economic interest and contractual rights. FATF guidance on beneficial ownership of legal persons and legal arrangements provides authoritative concepts and risk-oriented guidance [12-13]. UAE Cabinet Resolution No. 109 of 2023 addresses beneficial-owner procedures within its scope [11]. Qualified compliance and legal owners determine applicability and acceptance.

6.4 Document classification and completeness

Document classification should return document type, candidate entity, period, status, language, signature state, source authority, confidence and exceptions. A low-confidence classification routes to review. Completeness is tested against the applicable workplan. The system should state exactly which expected items are absent, stale, unsigned, partial or outside permission.

Intake failure	Example	Required route
unreadable source	scan quality blocks reliable extraction	request replacement or approved OCR review
permission ambiguity	deck forwarded without clear onward-sharing right	information owner and legal review
entity conflict	name matches two registry entities	compliance or legal resolution
period gap	model omits a required historical month	counterparty request and analyst review
version conflict	two files claim to be final	document owner resolution
embedded instruction	source text attempts to redirect the model	treat as untrusted content and follow system policy
executable content	file contains active code or macro	isolate and follow security procedure

6.5 Electronic records and evidence

UAE Federal Decree-Law No. 46 of 2021 addresses electronic transactions and trust services [16]. UAE Federal Decree-Law No. 35 of 2022 addresses evidence in civil and commercial transactions [17]. These sources can inform record design. They do not make every stored AI output legally admissible or prove authenticity. Qualified counsel determines evidential effect. The system should preserve original sources, hashes, timestamps, transformations, signatures and approval records.

7. ELIGIBILITY, MANDATE AND SUITABILITY SCREENING

7.1 Gate sequence

Screening should apply gates in an explicit order:

1. source and permission gate;
2. entity and restricted-party gate;
3. mandate hard-exclusion gate;
4. delegated-authority gate;
5. minimum-information gate;
6. conflict and independence gate;
7. portfolio and liquidity gate;
8. specialist-routing gate; and
9. human progression decision.

The model may populate candidate values. Deterministic rules evaluate accepted values. A reviewer accepts, corrects or leaves each material value unresolved. This separation prevents an extracted statement from becoming an approved mandate fact.

7.2 Private-fund screening

ILPA's Due Diligence Questionnaire 2.0 provides a structured private-fund diligence resource [23-24]. ILPA's ESG diligence materials provide additional guidance within their scope [25]. A family office can map these resources to its own approved questionnaire. The system should record which questions apply, the manager's response, supporting evidence, reviewer, exceptions and follow-up. A completed questionnaire is a representation set, not independent verification.

7.3 Direct-investment screening

Direct investments require a distinct schema: target entity, owners, management, business model, financial history, transaction structure, valuation, governance rights, funding plan, regulatory perimeter, legal risks, tax questions, ESG or reputational issues, exit pathways and monitoring. The system should not force a direct deal into a fund template.

7.4 Suitability and portfolio context

The portfolio layer should calculate exposure using approved, versioned definitions. It may show committed capital, funded capital, unfunded commitments, look-through exposure where supported, currency, sector, geography, liquidity buckets, counterparty concentration and scenario cash needs. Every aggregation must retain lineage to the underlying position and valuation date.

Suitability dimension	Evidence	Decision question
liquidity	cash, near-cash, commitments, calls, distributions and family needs	can the family meet obligations under approved scenarios?
concentration	current and proposed exposures under approved taxonomy	does the proposal breach or approach a limit?
currency	denomination, hedging, cash needs and policy	is currency exposure authorised and understood?
leverage	fund, vehicle, asset and family-level debt where available	which leverage layers affect loss and liquidity?
governance	board, veto, information and exit rights	are rights consistent with the family's control objectives?
complexity	entities, jurisdictions, valuation and operating dependencies	which specialist skills and monitoring are required?
reputation	supported adverse information and relationship effects	who determines acceptability and mitigation?

7.5 Diligence standard

CFA Institute Standard V(A) addresses diligence and reasonable basis for members and candidates subject to the standard [26]. A family office may adopt a comparable internal principle regardless of direct applicability. The system should support an adequate basis by linking conclusions to relevant evidence and disclosing limitations. It cannot determine that diligence is legally or professionally sufficient.

8. DILIGENCE EVIDENCE GRAPH

8.1 Atomic evidence object

A memo paragraph is too large to be the primary evidence unit. Each material proposition should be represented as an atomic evidence object. The object records subject, predicate, value, units, period, perimeter, source, exact location, source authority, evidence class, transformation, permission, reviewer, status and refresh trigger.

Evidence classes should remain explicit. **Observed** describes a value obtained from an approved source or measurement. **Represented** describes a statement by a counterparty or adviser. **Derived** describes a calculation or transformation. **Determined** describes a conclusion made by an authorised specialist or decision maker. A model output is a candidate derived object until validated and accepted.

8.2 Provenance model

W3C PROV-O supplies a standard vocabulary for entities, activities and agents [40]. The family-office evidence graph can extend that model with opportunity, mandate, organisation, legal entity, ownership relation, document, claim, table cell, calculation, workstream, issue, conflict, decision and monitoring obligation. Each relationship should retain time and version.

Bi-temporal history is useful. Valid time records when a fact applied in the investment world. System time records when the office learned, recorded or corrected it. This allows a later reviewer to reconstruct the information available at the original decision.

8.3 Evidence graph relationships

From	Relationship	To	Example control
opportunity	introduced by	person or organisation	preserve source and permission
legal entity	owned or controlled by	person, entity or arrangement	record source, percentage, role and date
document	supports or contradicts	claim	retain exact source location
calculation	uses	evidence object	preserve formula, units and input versions
claim	affects	mandate rule or analysis	show consequence and materiality
issue	blocks or conditions	stage transition	require owner, due date and closure evidence
memo sentence	cites	evidence or analysis record	sentence-level support test
decision	creates	condition or monitoring obligation	promote into the portfolio workflow

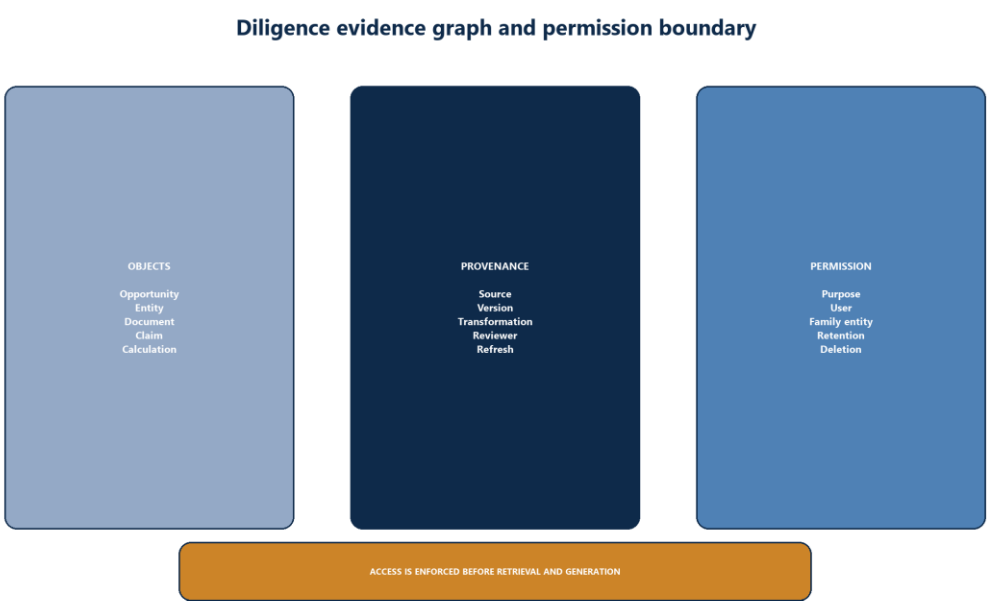

Figure 4. Diligence evidence graph and permission boundary

Source: Matchpoint evidence architecture; data, beneficial-ownership and provenance context [9-21,40].

8.4 Contradictions and absence

The graph must retain contradictions. If a deck reports one revenue figure and audited statements report another, neither value should be silently overwritten. A conflict object links both claims, states the comparison basis and routes resolution. A missing document is an unresolved evidence condition. It is not evidence that the underlying fact is false.

The system should use bounded language. “No litigation was identified in the supplied sources listed in the search record as at the stated date” is different from “the company has no litigation.” The second statement requires appropriate evidence and authority.

8.5 Evidence freshness

Every material object should have a refresh rule: fixed at transaction date, refreshed on document receipt, periodic, event-driven or expiring. Registry status, sanctions screening, cash, portfolio exposure, valuation and process deadlines change at different rates. A memo should show the evidence date and flag stale items before release.

9. DILIGENCE-EXTRACTION WORKFLOW

9.1 Task decomposition

Anthropic's engineering guidance distinguishes workflows with predefined code paths from agents that dynamically direct tool use [29]. Deal diligence benefits from controlled workflows for predictable tasks and tightly bounded agent behaviour for complex retrieval or comparison. A broad prompt such as "underwrite this deal" lacks a defined output, source boundary and authority. The operating system should decompose work into testable tasks.

Typical tasks include document classification, table extraction, entity reconciliation, contract-clause comparison, management-question tracking, financial normalisation, ownership mapping, evidence retrieval, issue drafting and memo assembly. Each task declares inputs, allowed sources, schema, validation, abstention condition, reviewer and downstream use.

9.2 Extraction contract

The extraction contract should require:

1. exact source location, including page, section, table or cell where available;
2. faithful source value and normalised value as separate fields;
3. units, currency, period and entity perimeter;
4. qualifiers, assumptions and footnotes;
5. null for absent data;
6. confidence used only for review routing;
7. validation result and exception list; and
8. reviewer disposition.

Anthropic's citation documentation describes product functionality for grounding generated responses in supplied sources [31]. Citation presence is one control. A separate support test should determine whether the cited passage supports the full claim, applies to the correct entity and period, and has adequate authority.

9.3 Structured and deterministic tools

Arithmetic, currency conversion, date logic, ownership aggregation, portfolio limits and valuation formulas should run in versioned deterministic tools. Natural-language components may select an approved tool, prepare inputs and explain outputs. The calculation record remains authoritative. Anthropic's advanced tool-use materials can inform implementation patterns [33]. Product capability does not validate a specific family-office deployment.

9.4 Diligence-extraction sequence

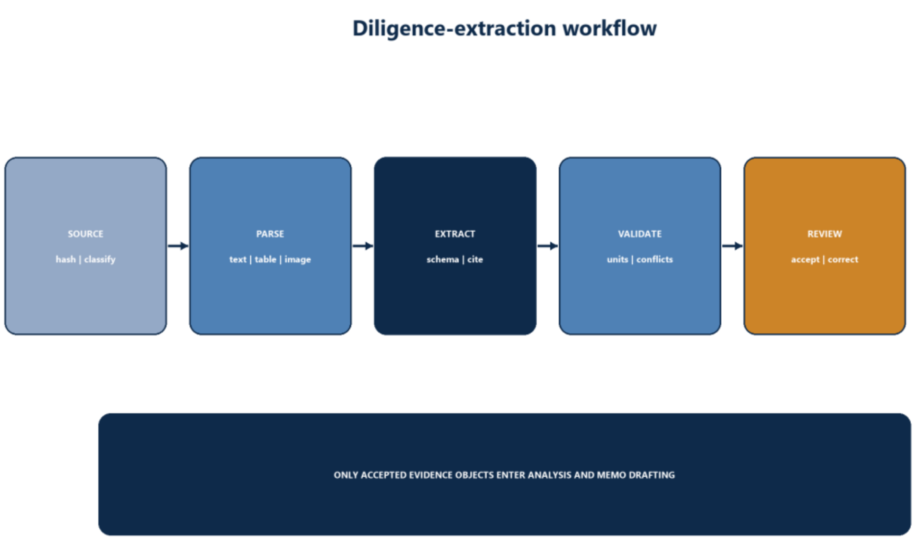

Figure 5. Diligence-extraction workflow
Source: Matchpoint extraction framework; Claude tool, citation and evaluation context [29-33].

The workflow sequence is:

1. preserve and hash the source;
2. scan and classify;
3. parse text, tables, images and metadata;
4. resolve candidate entities and periods;
5. extract schema-bound candidate facts;
6. validate types, units, totals and cross-references;
7. link candidates to exact source locations;
8. compare against existing evidence;
9. create conflicts, gaps and questions;
10. route material items to the named reviewer;
11. accept, correct, condition or reject; and
12. promote accepted objects into analysis and drafting.

9.5 Questions and follow-up

Management and adviser questions should be generated from unresolved objects, not generic templates alone. Each question links to the missing or conflicting evidence, states the requested form, identifies the decision affected and preserves the response. A response remains a representation until the review standard is met.

9.6 Language and translation

GCC transactions may involve Arabic and English sources. Machine translation can assist discovery and comparison. The system should retain original language, translation engine or provider, date, reviewer and purpose. Material legal, regulatory or contractual interpretation should be confirmed by appropriately qualified bilingual specialists.

10. ENTITY, BENEFICIAL-OWNERSHIP AND KYC EVIDENCE

10.1 Entity perimeter

The entity map should include the capital owner, investment vehicle, target, sellers, sponsor, manager, general partner, investment adviser, administrators, key subsidiaries, financing entities and material counterparties. Each entity should have jurisdiction, legal form, registration identifier, status, addresses, directors, authorised signatories and linked official evidence where available.

10.2 Ownership and control

FATF guidance describes beneficial-ownership transparency for legal persons and legal arrangements [12-13]. UAE Cabinet Resolution No. 109 of 2023 contains beneficial-owner procedures within its scope [11]. The workflow should represent direct and indirect interests, voting, appointment rights, control by other means, nominee arrangements and trust or foundation roles as distinct relationships. A percentage calculation should preserve every input, date and source.

10.3 AML and restricted-party controls

UAE Federal Decree-Law No. 10 of 2025 addresses anti-money-laundering, combating the financing of terrorism and proliferation financing and superseded the cited earlier federal decree-law [10]. Its application depends on the entity and activity. The family office's approved compliance owner and qualified advisers should specify screening populations, lists, frequency, escalation and record retention.

The language model may extract candidate names and identifiers. Approved screening systems and authorised reviewers perform and resolve checks. A name-only match should not be treated as identity. False positives, aliases, transliteration and incomplete dates of birth require controlled handling.

10.4 Ownership-evidence matrix

Object	Required evidence	Common exception	Decision owner
legal existence	current registry record or equivalent authoritative evidence	stale or inaccessible registry	compliance or legal
direct ownership	register, filing, constitutional record or transaction document	conflicting percentages	legal and compliance
indirect ownership	supported chain through each intermediate entity	missing intermediate record	compliance
control	voting, appointment, contractual or other supported control right	control claim without document	legal
authorised signatory	board resolution, power or official mandate	expired or limited authority	legal and operations
restricted-party result	approved system result with identifiers and date	potential match	compliance
source of funds or wealth	evidence under approved policy	incomplete or inconsistent evidence	compliance

10.5 Family confidentiality

Family structures can contain sensitive personal and commercial information. Access should be purpose-bound and role-based. Principals, family members, investment teams, advisers and service providers may require different views. Retrieval must enforce permission before content enters model context. Redaction after generation is insufficient as the sole control because the content may already have been processed outside the permitted boundary.

11. FINANCIAL, COMMERCIAL AND RISK WORKSTREAMS

11.1 Workstream matrix

The workplan should be generated from transaction type, stage, sector, geography, instrument and known risks, then approved by the deal lead. Common workstreams include commercial, financial, tax, legal, regulatory, technical, operational, cybersecurity, ESG, reputation, insurance, valuation, governance and portfolio fit. The system should store applicability, scope, owner, provider, source population, questions, findings, exceptions and sign-off.

11.2 Financial normalisation

Financial data should preserve entity, consolidation perimeter, accounting basis, currency, units, period, source status and audit status. Every adjustment should record source, rationale, amount, sign, tax or cash treatment where relevant, recurrence assumption, scenario, reviewer and version. The model may propose a candidate normalisation. The authorised analyst accepts or changes it.

Financial test	Evidence question	Control
revenue	which entity, period, recognition basis and source?	reconcile statements, ledger and operating data where supplied
margin	which cost classifications and adjustments apply?	retain reported and adjusted views
cash conversion	how do earnings reconcile to cash?	deterministic bridge with working-capital lineage
debt	which instruments, guarantees and off-balance commitments exist?	legal and financial reconciliation
forecasts	which assumptions are represented, contracted or independently supported?	separate cases and record owner
valuation	which method, inputs, date and sensitivity apply?	versioned calculation and authorised review

11.3 Commercial evidence

Commercial diligence should separate facts, representations, observations and projections. Customer concentration, pipeline, churn, pricing, unit economics, market size and competitive position require defined populations and dates. A management deck's market statement is a representation until supported by an appropriate source. Interview notes should identify participant, date, questions, consent and reviewer.

11.4 Risk taxonomy

The risk register should link each risk to cause, affected objective, evidence, likelihood definition, impact definition, controls, mitigation, owner, residual assessment, trigger and monitoring obligation. Scores may support prioritisation only when definitions and scales are approved. A single aggregate number can hide a severe tail risk or an unresolved legal question.

11.5 Portfolio and liquidity analysis

The investment should be analysed with current portfolio data and approved scenarios. For a fund commitment, the system may model calls, distributions and unfunded exposure using stated assumptions. For a direct investment, it may model funding tranches, follow-on needs, dilution and exit timing. All scenarios should be labelled. Forecasts are not observed facts.

11.6 Specialist evidence

External reports should be ingested with provider, engagement scope, reliance terms, date, version and permission. A red-flag report, full-scope diligence report and legal opinion have different evidential roles. The system should not extend a specialist conclusion beyond its stated scope or recipient.

12. INVESTMENT-COMMITTEE MEMO GENERATION

12.1 Memo as a compiled decision record

The IC memo should be compiled from approved evidence and analysis records. It is not an unconstrained summary of the data room. Each material sentence links to one or more sources, calculations, specialist determinations or named judgements. The memo states the evidence date, mandate version, portfolio snapshot and unresolved items.

A reference structure is:

1. decision requested and authority;
2. executive judgement by named investment professional;
3. opportunity and transaction overview;
4. mandate and portfolio fit;
5. investment case and evidence;
6. principal risks and mitigants;
7. valuation, returns and sensitivities;
8. diligence status and specialist findings;
9. conflicts and governance;
10. conditions, monitoring and exit;
11. recommendation and alternatives; and
12. source, model and approval appendices.

12.2 Claim ledger

Before drafting prose, the workflow should create a claim ledger. Each row states claim text, class, supporting object, source location, date, reviewer, materiality, permitted wording and memo destination. Unsupported material claims stay out of the memo or appear as explicit unresolved issues.

Claim class	Permitted wording	Required support
observed fact	direct statement with date and perimeter	accepted evidence object
counterparty representation	“management represents” or equivalent	dated source and representative
derived analysis	“the approved model calculates”	formula, inputs, version and reviewer
specialist determination	attributed to named role or provider within scope	report, date, scope and reliance status
investment judgement	attributed to named author or committee	evidence basis and decision record
unresolved matter	explicit uncertainty and consequence	issue record, owner and due date

12.3 Drafting sequence

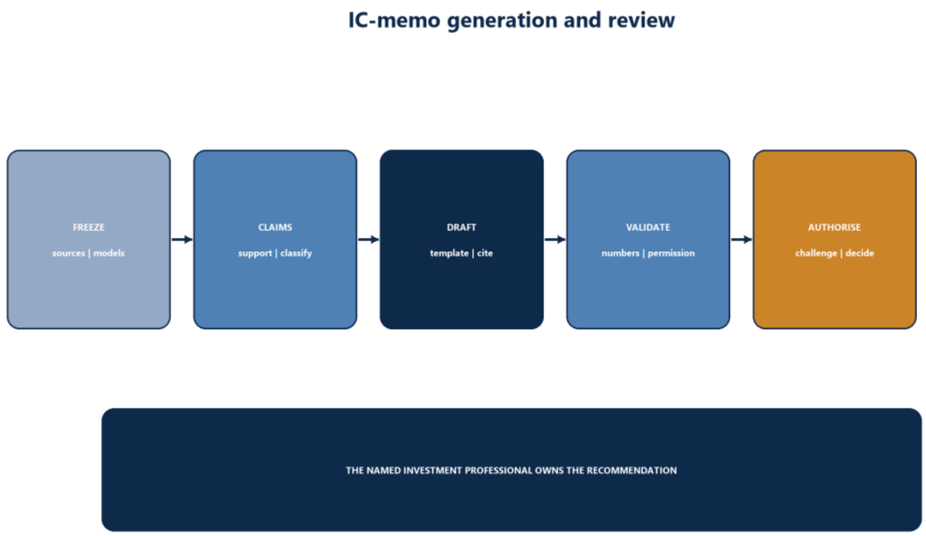

Figure 6. IC-memo generation and review sequence

Source: Matchpoint claim-ledger and committee-control framework [26-33,37-40].

The system first freezes the source population and calculation versions. It assembles the claim ledger, tables and exception schedule. It drafts section-level prose within a controlled template. A validation pass checks citation support, numeric agreement, terminology, permission, contradictions and missing sections. The named investment professional edits the judgement, recommendation and emphasis. Specialists review their areas. The authorised decision body receives a versioned pack.

12.4 Numerical integrity

Every number in the memo should originate from an accepted source or approved calculation. Units and currencies should appear in the data model. Repeated numbers should link to one canonical object. A changed model input should invalidate affected sentences and tables until regenerated and reviewed.

12.5 Citation support

Citation support should be tested at claim level. A cited page can contain the same topic without supporting the stated conclusion. The checker should test entailment within bounded criteria, then route material results to a reviewer. Anthropic's evaluation guidance supports defining success criteria and empirical tests [32]. A vendor document does not replace independent validation.

12.6 Recommendation authorship

The recommendation should identify the human author and authority. Model-generated wording can assist structure and consistency. The system should not present a model as the investment decision maker. Model identity, configuration, source population and validation results remain in the run record rather than being concealed.

13. HUMAN DECISION RIGHTS, CONFLICTS AND CONTROL

13.1 Segregation of duties for a two-person team

A compact team cannot always create institutional headcount-based segregation. It can create transaction-level controls. The proposer and reviewer should be named. The reviewer should challenge evidence, calculations, conflicts and recommendation. Material legal, tax, compliance, valuation and technical matters route to qualified external or internal specialists. The principal or committee retains approval where required.

Where one person performs multiple steps, the system should disclose the overlap and require compensating review defined by policy. A model cannot serve as independent reviewer of its own work.

13.2 Approval matrix

Action	Preparer	Reviewer	Approver
opportunity record	investment professional	relationship or operations owner	delegated owner if required
mandate result	system plus investment professional	CIO or policy owner	authorised owner for exceptions
entity and ownership acceptance	analyst or provider	compliance and legal as applicable	authorised onboarding owner
valuation model	analyst or adviser	independent qualified reviewer	authorised investment authority
specialist conclusion	specialist	deal lead within reliance scope	decision body considers it
memo judgement	named investment professional	challenge reviewer	principal or committee
release of communication	drafter	information owner and specialists as needed	authorised sender

13.3 Override governance

Overrides are permitted only through an approved route. Each override records the rule, original result, changed result, reason, evidence, authority, date, expiry and monitoring. A high override rate may indicate poor rules, weak intake or inappropriate pressure. It is a monitoring signal, not automatic proof of misconduct.

13.4 Committee record

The committee record should preserve attendees, authority, materials, conflicts, questions, decisions, conditions, dissents, abstentions and follow-up. Generated minutes are drafts until approved. Conditions should become structured obligations with owner, due date and evidence requirement.

13.5 Communications

Any message to a sponsor, manager, co-investor, adviser or family member should use only approved information and the correct permission context. Legal or tax communications route through qualified advisers where appropriate. The system should preserve the approved version, sender and recipients.

14. WORKED DEAL-SCREENING SCENARIO

14.1 Scenario status

The following case is an **[Unverified illustrative scenario]**. It does not describe a Matchpoint client, family office, transaction or observed result. Names, amounts, team capacity, times, costs, scores and outcomes are invented solely to demonstrate the operating architecture. No investment conclusion follows from the example.

14.2 Opportunity

An unnamed GCC family office receives a co-investment proposal involving a private operating company. The proposal arrives through an adviser and includes a teaser, management presentation, three-year financial workbook, draft term sheet, ownership chart and data-room link. The investment team has two members. The family mandate, committee authority and current portfolio snapshot are assumed to exist as approved system inputs.

The system preserves the source files, creates an opportunity, links the introduction and tests permission. It extracts candidate entity names, transaction type, proposed investment range, currency, sector, geography, process date and adviser. The team confirms or corrects those fields.

14.3 Initial screening

The hard-gate engine identifies no assumed prohibited category under the fictional mandate. It flags three missing fields: final target registration number, complete beneficial-ownership evidence and a confirmed decision timetable. The portfolio layer shows a candidate sector concentration close to an invented review threshold. The system therefore proposes a conditioned progression. The CIO determines whether the case enters diligence.

14.4 Evidence extraction

The financial workbook contains three currencies across tabs and inconsistent units. The extraction workflow creates separate candidate objects, validates totals and flags a discrepancy between the deck and workbook revenue. The ownership chart names an intermediate holding company without a registry identifier. The draft term sheet includes a governance right that differs from the teaser. Each exception links to the exact source location and an owner.

14.5 Workplan and specialist routing

The team selects approved commercial, financial, legal, tax, compliance and cybersecurity workstreams. The system generates a requirements matrix. Beneficial-ownership and transaction-structure questions route to qualified compliance and legal advisers. The model does not resolve them.

14.6 IC memo

The memo compiler uses accepted evidence, an approved financial model and specialist records. It attributes management representations, cites source locations and lists unresolved conditions. The named investment professional writes the recommendation. The challenge reviewer tests the downside assumptions, concentration treatment and conditions.

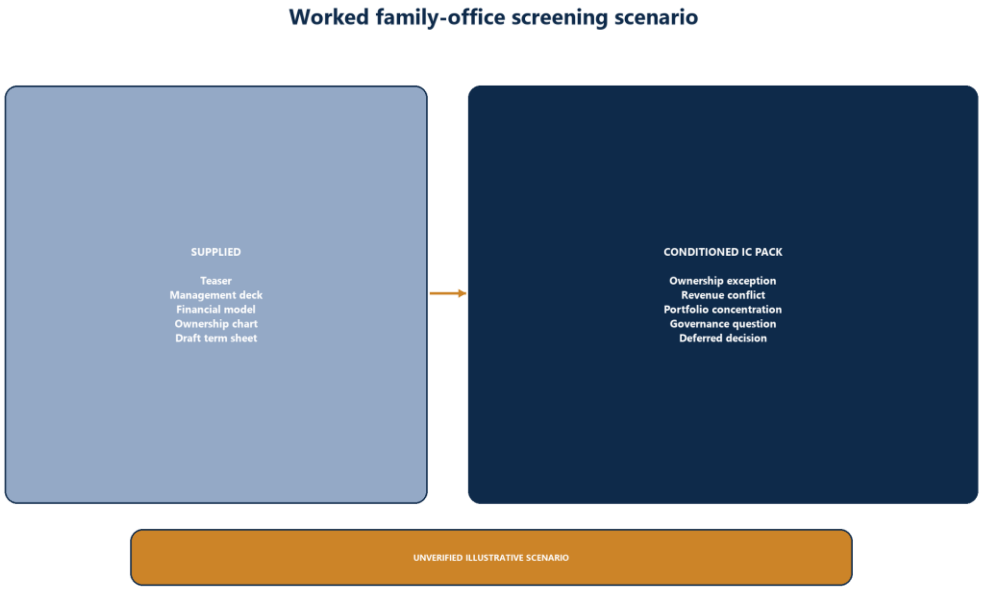

Figure 7. Worked scenario from source intake to conditioned IC pack

Source: Matchpoint framework. All names, values and events are unverified illustrative assumptions.

14.7 Outcome

The fictional committee defers the decision pending ownership evidence, reconciliation of the revenue discrepancy and clarification of governance rights. The system converts those items into conditions with owners and due dates. No invented speed, headcount, cost or investment-return benefit is claimed.

15. GOLDEN-SET EVALUATION

15.1 Evaluation before production

Anthropic's evaluation guidance recommends defining success criteria and empirical tests [32]. NIST's AI Risk Management Framework and AI Resource Center provide voluntary risk-management resources [37-38]. A family-office deployment should translate those materials into task-specific tests, named owners and release thresholds.

The golden set should represent the office's actual opportunity types, languages, document formats, jurisdictions and failure modes, subject to permission. Cases should contain authoritative answers or approved reviewer labels. The evaluation population should remain separate from prompt and workflow development.

15.2 Component tests

Component	Primary measure	Material failure example
document classification	precision and recall by class and status	executed agreement labelled draft
entity resolution	pairwise precision and recall	two distinct entities merged
table extraction	cell and row accuracy with units	currency or sign error
ownership mapping	supported relationship accuracy	controller omitted or unsupported
mandate gating	agreement with approved deterministic result	hard exclusion missed
retrieval	authoritative evidence recall within permission	material source omitted or unauthorised source exposed
citation support	complete-claim support and exact location	citation does not support conclusion
numeric integrity	agreement with approved calculations	return, exposure or valuation error
issue detection	severity-weighted recall	material conflict or missing evidence not flagged
memo drafting	factuality, completeness, attribution and structure	representation presented as fact
workflow	correct state, owner and approval route	model bypasses required reviewer

15.3 End-to-end tests

End-to-end cases should ask whether the workflow reaches the correct controlled state. Test cases should include supported progression, conditional progression, rejection, deferment, security exceptions, privilege restrictions, stale evidence, inconsistent currency, ownership ambiguity, conflicting financials, hidden prompt injection, tool failure and model unavailability.

15.4 Severity weighting

Errors should be weighted by consequence. A formatting defect differs from unauthorised disclosure, a missed hard exclusion, a currency error or a false beneficial-owner conclusion. Release criteria should combine aggregate performance with zero-tolerance conditions for defined critical failures. A strong average cannot offset a critical security or decision-rights failure.

15.5 Blinded human review

Where practical, reviewers should assess outputs without knowing whether the initial draft was human or AI-assisted. The study should use the same evidence population, templates and decision gates. Reviewer agreement and adjudication should be recorded. A model-as-judge score can support development; it cannot be the only material acceptance test.

15.6 Production monitoring

Monitoring should track task volume, input drift, exception rates, abstention, reviewer corrections, citation failures, numerical failures, permission incidents, overrides, cycle times, vendor changes and model changes. A version change should trigger regression testing. Production incidents should have containment, notification, fallback, root-cause review and approved restart procedures.

16. SECURITY, PRIVACY AND VENDOR GOVERNANCE

16.1 Data map and classification

The family office should map personal data, family information, portfolio data, trade secrets, privileged material, restricted lists, credentials and system logs. Each class should have an owner, purpose, permitted users, location, retention, deletion and incident route. UAE, DIFC and ADGM data-protection requirements apply according to their respective scopes [9,14-15]. Qualified privacy and legal advisers determine the applicable obligations.

16.2 Access before retrieval

Identity, role, purpose, opportunity, family entity and source permission should be evaluated before retrieval. The retrieval service should return only authorised objects and record what it returned. Model context, tool calls, caches and logs require the same policy. Post-generation redaction can provide an additional control. It does not remove the need for permissioned retrieval.

16.3 Prompt injection and untrusted content

Data-room documents, websites, email and attachments are untrusted content. Instructions embedded in them should be treated as data. The model follows the approved system and workflow authority. Tools should validate parameters and restrict actions to the minimum approved scope. High-impact actions require deterministic gates and human approval.

NIST AI 100-2 provides a taxonomy and terminology for adversarial machine-learning attacks and mitigations [39]. It can support threat modelling. It does not certify a deployment.

16.4 Vendor and model record

Central Bank of the UAE enabling-technologies and outsourcing materials apply to licensed financial institutions within their scope [18-19]. A family office outside that perimeter may use the control concepts as benchmarks while recording that status. The vendor file should cover service description, data use, training policy, hosting, subprocessors, encryption, access, retention, deletion, audit evidence, incident notification, business continuity, exit, intellectual property and change management.

Anthropic's trustworthy-agent research and autonomy measurements can inform control design [30,36]. Each implementation still requires independent evaluation. The production run record should identify provider, model, configuration, system instructions, task prompt, tools, retrieval policy, source population, validations, reviewer and downstream use.

16.5 Incident and fallback

The operating plan should cover unauthorised access, data leakage, supplier outage, source corruption, prompt injection, incorrect calculation, false citation, model change and workflow failure. Fallback may mean manual processing, an approved secondary provider or postponing the decision. A service-level target should not force release during a control failure.

17. COST AND CAPACITY MEASUREMENT

17.1 Evidence boundary

This paper contains no approved observed Matchpoint or client productivity, cost, risk or investment-return evidence. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0. The cost and capacity framework below uses formulas and **[Unverified illustrative scenarios]** only. Any local inputs require CK approval and finance evidence before a public claim.

17.2 Cost model

The fully loaded annual operating cost should include approved cash compensation, employer costs, benefits, workspace, software, data, advisers attributable to the operating model, implementation, security, support, model usage and governance. The comparison should avoid treating existing sunk costs as incremental or excluding retained reviewer and specialist effort.

Cost component	Formula	Required evidence
internal labour	approved loaded hourly rate multiplied by observed active hours	payroll or finance-approved rate and time record
technology	licence, usage, infrastructure and support	contract, invoice and allocation method
implementation	approved internal and external project cost	invoices and approved internal-cost method
specialist review	incremental legal, tax, compliance, valuation and technical cost	engagement and invoice evidence
control operation	security, evaluation, monitoring, audit and incident readiness	approved operating budget
rework	loaded rate multiplied by observed correction time	version and review records
incident cost	approved direct and indirect event cost	incident and finance record

17.3 Capacity model

Decision capacity should be modelled by case complexity and bottleneck role. Let Q be quality-gated decision-ready packs, H be approved active human hours, A be available hours, and U be utilisation permitted by policy. Observed productivity can be stated as Q divided by H. Sustainable capacity should also consider adviser constraints, committee cadence, portfolio-monitoring duties and peak workload. A model-derived capacity estimate is not an observed result.

17.4 Two-person versus ten-person protocol

The parity test should preregister:

1. team composition and roles;
2. comparator role composition and specialist access;
3. case population and complexity strata;
4. source completeness and languages;
5. common mandate, templates and quality gates;
6. measurement period and warm-up;
7. permitted tools and training;
8. primary and secondary outcomes;
9. failure and incident definitions;
10. reviewer blinding and adjudication;
11. statistical or decision rule; and
12. limits on generalisation.

Cases should be randomly assigned where feasible or closely matched. The result should report uncertainty and the exact evaluated scope. If two people produce the same count with more severe errors, more adviser work or unsustainable hours, parity has not been established.

17.5 Worked formula

Assume only for illustration that a baseline and pilot use comparable cases. Define baseline cost per decision-ready pack as total approved operating cost for the measured population divided by packs passing all gates. Define pilot cost on the same basis. Net operating value is approved baseline cost minus approved pilot cost, less implementation and transition costs, with any quality, risk or investment outcome claimed separately from observed evidence. No values are inserted because none have been approved for T34.

Figure 8. Cost, capacity and quality measurement bridge

Source: Matchpoint measurement framework; external-domain evidence [41-43]. No family-office benefit is observed.

17.6 Public-claim gate

A public productivity or cost claim requires an approved protocol, comparable population, stable workflow, observed period, quality results, failures, cost evidence, reviewer approval and legal or communications review where required. Cross-domain studies may be cited as context [41-43]. They cannot be converted into a family-office headcount claim.

18. NINETY-DAY IMPLEMENTATION ROADMAP

18.1 Days 0 to 15: scope and authority

Select one deal type, one family capital entity and one decision route. Name the principal or policy owner, CIO, deal lead, challenge reviewer, information owner, privacy owner, compliance owner, security owner and specialist routes. Approve the mandate version, decision states, minimum evidence and release authority.

18.2 Days 16 to 30: source and evidence foundation

Create the source inventory, data classes, permission model, opportunity object, entity schema, evidence object, claim ledger and exception record. Configure immutable originals, file hashes, audit history, retention and deletion. Build a small golden set from permissioned historical cases.

18.3 Days 31 to 45: screening and extraction

Implement intake classification, entity matching, mandate hard gates, requirements matrix, schema-bound extraction and deterministic validation. Test prompt injection, inaccessible sources, missing fields, conflicting units and document versions. Keep every output in reviewer mode.

18.4 Days 46 to 60: diligence and memo shadow mode

Add approved workstream templates, evidence graph, retrieval, claim ledger and IC-memo drafting. Run cases in shadow mode alongside the current process. Review every material claim, citation, number, issue and state transition. Strengthen the golden set with observed failures.

18.5 Days 61 to 75: controlled pilot

Run a bounded pilot with approved cases and no autonomous consequential action. Measure active time, cycle time, review effort, corrections, material failures, specialist work, model use and cost. Hold release if a critical control fails.

18.6 Days 76 to 90: governance decision

Compare pilot evidence with preregistered gates. Review security, privacy, vendor, quality, incident and fallback evidence. The authorised family-governance body decides whether to stop, revise, extend shadow mode or approve a limited production scope. Production permissions should be narrower than the successful test boundary where risk warrants.

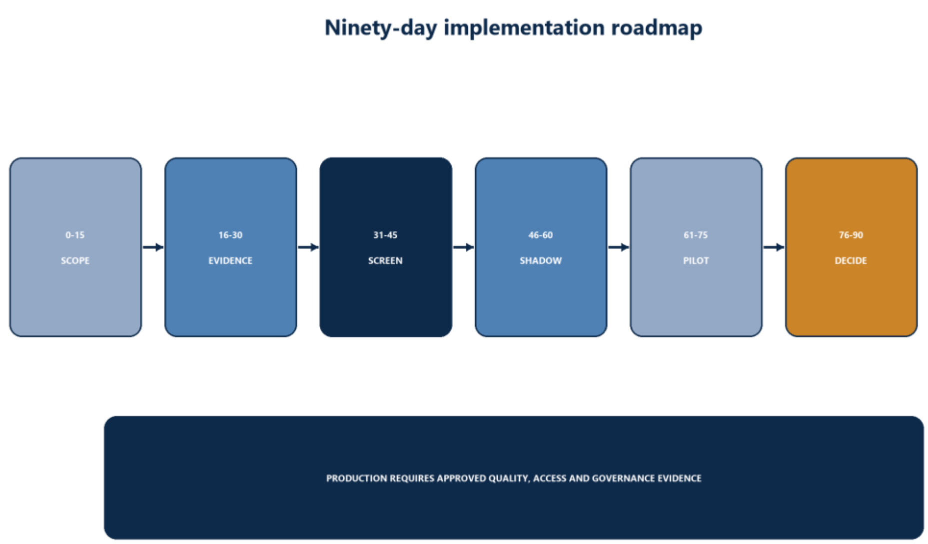

Figure 9. Ninety-day implementation roadmap

Source: Matchpoint implementation framework; governance context [29-40].

18.7 Minimum deliverables

The minimum operating pack includes the approved mandate configuration; decision-rights matrix; source and permission inventory; opportunity, entity, evidence, exception, claim and run schemas; requirements matrices; deterministic rule and calculation register; evaluation set and results; security and privacy assessment; vendor file; incident and fallback plan; and observed-benefits register.

19. RELEASE CHECKLIST AND CLAIMS REGISTER

19.1 Decision-pack release checklist

- opportunity, source, sponsor, target, vehicle and material entities are identified or explicitly unresolved;
- source permission, confidentiality and permitted recipients are confirmed;
- applicable mandate version and hard gates are recorded;
- portfolio snapshot, liquidity and concentration definitions are current;
- beneficial-ownership and restricted-party processes have completed under approved authority;
- financial periods, currencies, units, perimeter and adjustments reconcile;
- valuation method, inputs, date, scenarios and reviewer are recorded;
- workstreams are complete or carry approved exceptions;
- specialist conclusions stay within scope and reliance terms;
- every material memo claim has accepted support or explicit uncertainty;
- every number agrees with its approved source or calculation;
- conflicts, overrides, conditions and dissent are disclosed;
- the recommendation identifies its human author;
- authorised reviewers have approved the pack; and
- approved conditions and monitoring obligations are structured for follow-up.

19.2 Public claims register

Claim	Evidence status in T34	Permitted wording
Claude can support schema-bound extraction, comparison, citation and drafting	supported as product or architecture capability within cited sources [29-33]	state capability with controlled-deployment qualification
a controlled evidence graph can preserve source and transformation lineage	supported as an operating architecture using provenance concepts [40]	describe design, not observed benefit
two people can match a ten-person institutional team	unsupported	identify as unverified hypothesis; do not claim
AI reduces family-office diligence cost	unsupported	do not claim
AI improves family-office investment outcomes	unsupported	do not claim
cross-domain studies report task productivity effects	supported within their study scopes [41-43]	cite exact domain and limitations
Matchpoint or client revenue, saving, loss reduction or alpha resulted	no approved observed evidence	record USD 0 and do not claim benefit

19.3 Model-output language

Permitted labels include candidate extraction, represented fact, supported evidence object, derived calculation, specialist determination, investment judgement, conditioned item and unresolved item. Prohibited language includes “guaranteed”, “risk-free”, “fully autonomous”, “legally compliant” or “investment-ready” unless an appropriately authorised determination supports the exact statement and scope.

19.4 Governance release

Production approval should identify eligible users, opportunities, jurisdictions, models, tools, permissions, providers, tasks, thresholds and expiry. It should state which actions always require human approval. Any model, provider, mandate or material workflow change should trigger the specified review and regression tests.

20. LIMITATIONS AND CONCLUSION

20.1 Limitations

This is an operating-model and control paper. It is not a legal, regulatory, tax, compliance, accounting, audit, valuation, investment, privacy, cybersecurity or technology opinion. The cited laws, regulations, guidance and standards have distinct jurisdictions, scopes, statuses and effective dates. Qualified specialists should confirm current applicability and effect.

The paper does not contain a controlled family-office productivity experiment. The two-person-versus-ten-person proposition remains unverified. External studies concern customer support, consulting or software development and should not be generalised to investment diligence [41-43]. Anthropic product and research sources describe particular systems, observations or methods; they do not validate this architecture in production [29-36].

Language models can produce unsupported, incomplete or misleading outputs. Citations can be present while failing to support a full claim. Extraction can lose units, qualifiers or entity perimeter. Retrieval can omit a material source. Tools can fail. Vendors and models can change. The control design reduces or detects specified risks only to the degree demonstrated by evaluation and operations. It does not eliminate risk.

The worked scenario contains invented facts and no performance result. Cost and capacity formulas require local approved inputs. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence exists.

20.2 Conclusion

An AI-native family office should be designed as a decision-evidence system. The useful unit is the reviewable decision pack: sources preserved, permissions enforced, entities resolved, mandate gates visible, diligence structured, calculations reproducible, exceptions owned, claims cited and human authority recorded. Claude can support extraction, comparison, retrieval and drafting inside that architecture.

The two-person team hypothesis becomes meaningful only through a like-for-like, quality-gated pilot. The office should measure decision-ready throughput, human effort, rework, failures, specialist cost and sustainability. It should publish no headcount, cost or investment-result claim before observed evidence and approval.

The recommended starting point is one opportunity class and one decision route. Build the evidence objects, deterministic gates, claim ledger, evaluation set and release controls. Run shadow cases. Let the authorised family-governance body decide whether the evidence supports a limited production scope. That path can improve operating discipline even when the ultimate headcount hypothesis is rejected.

REFERENCES

[1] United Arab Emirates. Federal Decree-Law No. 37 of 2022 Concerning Family Companies. UAE Legislation. https://uaelegislation.gov.ae/en/legislations/1608. Accessed 2 August 2026.
[2] Government of Dubai. Decree No. 45 of 2022 Establishing the Family Companies Centre. Dubai Legislation Portal. https://dlp.dubai.gov.ae/Legislation%20Reference/2022/Decree%20No.%20%2845%29%20of%202022%20Establishing%20the%20Family%20Companies%20Centre.html. Accessed 2 August 2026.
[3] Government of Dubai. Law No. 9 of 2020 Regulating Family Property in the Emirate of Dubai. Dubai Legislation Portal. https://dlp.dubai.gov.ae/Legislation%20Reference/2020/Law%20No.%20%289%29%20of%202020%20Regulating%20Family%20Property%20in%20the%20Emirate%20of%20Dubai.html. Accessed 2 August 2026.
[4] Dubai International Financial Centre. DIFC Announces Enactment of New DIFC Family Arrangements Regulations. https://www.difc.com/whats-on/news/difc-announces-enactment-new-difc-family-arrangements-regulations. Accessed 2 August 2026.
[5] Abu Dhabi Global Market. Family Offices. https://www.adgm.com/business-areas/family-offices. Accessed 2 August 2026.
[6] Abu Dhabi Global Market. Foundations Regulations 2017. https://assets.adgm.com/download/assets/foundations-regulations-2017.pdf/a9b58092643811efb33122e97052245a?hash=6AFBC6314A0144F1577E66CBEB91D7E3&la=en. Accessed 2 August 2026.
[7] Abu Dhabi Global Market. ADGM Registration Authority Publishes Amendments to the Commercial Legislation. https://www.adgm.com/media/announcements/adgm-registration-authority-publishes-amendments-to-the-commercial-legislation. Accessed 2 August 2026.
[8] Abu Dhabi Global Market. Guidance and Policy Statements. https://www.adgm.com/legal-framework/guidance-and-policy-statements. Accessed 2 August 2026.
[9] Abu Dhabi Global Market. Data Protection Regulations 2021 Guidance, Part 1. https://assets.adgm.com/download/assets/ADGM%2BDPR%2B2021%2BGuidance%2BPart%2B1.pdf/b65534b2595411ef82c5a27efcbde115. Accessed 2 August 2026.
[10] United Arab Emirates. Federal Decree-Law No. 10 of 2025 on Anti-Money Laundering and Combating the Financing of Terrorism and Proliferation Financing. UAE Legislation. https://uaelegislation.gov.ae/en/legislations/3314. Accessed 2 August 2026.
[11] United Arab Emirates. Cabinet Resolution No. 109 of 2023 Concerning the Regulation of Beneficial Owner Procedures. UAE Legislation. https://uaelegislation.gov.ae/en/legislations/2176. Accessed 2 August 2026.
[12] Financial Action Task Force. Guidance on Beneficial Ownership of Legal Persons. March 2023. https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Guidance-Beneficial-Ownership-Legal-Persons.pdf.coredownload.pdf. Accessed 2 August 2026.
[13] Financial Action Task Force. Guidance on Beneficial Ownership and Transparency of Legal Arrangements. March 2024. https://www.fatf-gafi.org/content/dam/fatf-gafi/recommendations/Guidance-Beneficial-Ownership-Transparency-Legal-Arrangements.pdf. Accessed 2 August 2026.
[14] United Arab Emirates. Federal Decree-Law No. 45 of 2021 Concerning the Protection of Personal Data. UAE Legislation. https://uaelegislation.gov.ae/en/legislations/1972. Accessed 2 August 2026.
[15] Dubai Financial Services Authority. Data Protection. https://www.dfsa.ae/Data-Protection. Accessed 2 August 2026.
[16] United Arab Emirates. Federal Decree-Law No. 46 of 2021 on Electronic Transactions and Trust Services. UAE Legislation. https://uaelegislation.gov.ae/en/legislations/1539. Accessed 2 August 2026.
[17] United Arab Emirates. Federal Decree-Law No. 35 of 2022 Promulgating the Law of Evidence in Civil and Commercial Transactions. UAE Legislation. https://uaelegislation.gov.ae/en/legislations/1612. Accessed 2 August 2026.
[18] Central Bank of the UAE. Guidelines for Financial Institutions Adopting Enabling Technologies. https://rulebook.centralbank.ae/sites/default/files/en_net_file_store/CBUAE_EN_2413_VER1.pdf. Accessed 2 August 2026.
[19] Central Bank of the UAE. Outsourcing Regulation and Standards. https://rulebook.centralbank.ae/en/entiresection/2327. Accessed 2 August 2026.
[20] Dubai Financial Services Authority. Laws and Rules. https://www.dfsa.ae/your-resources/regulatory/laws-and-rules. Accessed 2 August 2026.
[21] Abu Dhabi Global Market. Legal Framework. https://www.adgm.com/legal-framework. Accessed 2 August 2026.
[22] United States Securities and Exchange Commission. Staff Responses to Questions About the Family Office Rule. https://www.sec.gov/rules-regulations/staff-guidance/division-investment-management-frequently-asked-questions/staff-responses-questions-about-family-office-rule. Accessed 2 August 2026.
[23] Institutional Limited Partners Association. Due Diligence Questionnaire. https://ilpa.org/resources-tools/resource-library/due-diligence-questionnaire/. Accessed 2 August 2026.
[24] Institutional Limited Partners Association. ILPA Due Diligence Questionnaire 2.0. https://ilpa.org/wp-content/uploads/2021/11/ILPA-DDQ-2.0.pdf. Accessed 2 August 2026.
[25] Institutional Limited Partners Association. Due Diligence and Investment Decision-Making. https://ilpa.org/industry-guidance/environmental-social-governance/esg-roadmap/due-diligence-and-investment-decision-making/. Accessed 2 August 2026.
[26] CFA Institute. Standard V(A): Diligence and Reasonable Basis. https://www.cfainstitute.org/standards/professionals/code-ethics-standards/standards-of-practice-v-a. Accessed 2 August 2026.
[27] CFA Institute. Standard III(C): Suitability. https://www.cfainstitute.org/standards/professionals/code-ethics-standards/standards-of-practice-iii-c. Accessed 2 August 2026.
[28] CFA Institute. Standard III(A): Loyalty, Prudence, and Care. https://www.cfainstitute.org/standards/professionals/code-ethics-standards/standards-of-practice-iii-a. Accessed 2 August 2026.
[29] Anthropic. Building Effective Agents. https://www.anthropic.com/engineering/building-effective-agents. Accessed 2 August 2026.
[30] Anthropic. Building and Evaluating Trustworthy Agents. https://www.anthropic.com/research/trustworthy-agents. Accessed 2 August 2026.
[31] Anthropic. Citations. Claude Documentation. https://docs.anthropic.com/en/docs/build-with-claude/citations. Accessed 2 August 2026.
[32] Anthropic. Define Success Criteria and Build Evaluations. Claude Documentation. https://docs.anthropic.com/en/docs/test-and-evaluate/define-success. Accessed 2 August 2026.
[33] Anthropic. Advanced Tool Use. https://www.anthropic.com/engineering/advanced-tool-use. Accessed 2 August 2026.
[34] Anthropic. Anthropic Economic Index Report: January 2026. https://www.anthropic.com/research/anthropic-economic-index-january-2026-report. Accessed 2 August 2026.
[35] Anthropic. Claude Code and Software-Development Expertise. https://www.anthropic.com/research/claude-code-expertise. Accessed 2 August 2026.
[36] Anthropic. Measuring Agent Autonomy. https://www.anthropic.com/research/measuring-agent-autonomy. Accessed 2 August 2026.
[37] National Institute of Standards and Technology. AI Risk Management Framework. https://www.nist.gov/itl/ai-risk-management-framework. Accessed 2 August 2026.
[38] National Institute of Standards and Technology. AI Resource Center. https://airc.nist.gov/. Accessed 2 August 2026.
[39] National Institute of Standards and Technology. Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, NIST AI 100-2 E2025. https://csrc.nist.gov/pubs/ai/100/2/e2025/final. Accessed 2 August 2026.
[40] World Wide Web Consortium. PROV-O: The PROV Ontology. https://www.w3.org/TR/prov-o/. Accessed 2 August 2026.
[41] Erik Brynjolfsson, Danielle Li and Lindsey R. Raymond. Generative AI at Work. NBER Working Paper 31161. https://www.nber.org/papers/w31161. Accessed 2 August 2026.
[42] Fabrizio Dell'Acqua et al. Navigating the Jagged Technological Frontier. Harvard Business School Digital Data Design Institute. https://aiinstitute.hbs.edu/navigating-the-jagged-technological-frontier/. Accessed 2 August 2026.
[43] Sida Peng et al. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. Microsoft Research. https://www.microsoft.com/en-us/research/publication/the-impact-of-ai-on-developer-productivity-evidence-from-github-copilot/. Accessed 2 August 2026.

APPENDIX A. OPPORTUNITY OBJECT

Field	Entry
opportunity ID	stable internal identifier
capital owner	family entity or authorised investment vehicle
source	person, organisation, channel and introduction record
permission	purpose, users, confidentiality, retention and deletion
sponsor or manager	linked entity identifier and role
target or vehicle	linked entity identifier and legal form
transaction	instrument, stage, amount range, currency and structure
classification	sector, geography, asset class and strategy under approved taxonomy
process	received date, deadlines, requested decision and stage
mandate result	policy version, rule results, reviewer and exceptions
owner	deal lead, challenge reviewer and specialist routes
status	controlled funnel state, date, reason and authority

APPENDIX B. EVIDENCE OBJECT

Field	Entry
evidence ID	stable, immutable identifier
subject	opportunity, entity, position, document or analysis object
predicate and value	atomic proposition and value
units and period	currency, scale, date or measurement interval
perimeter	entity, portfolio, instrument, scenario and exclusions
evidence class	observed, represented, derived or determined
source	document ID, version and exact location
authority	issuer, status, scope and reliance terms
transformation	tool, formula, input versions and output version
permission	permitted users, purpose, retention and deletion
validation	rules, results, exceptions and reviewer
status	candidate, supported, conditioned, conflicted or rejected
refresh	expiry, event trigger, owner and next review

APPENDIX C. IC MEMO CLAIM OBJECT

Field	Entry
claim ID	stable memo-specific identifier
memo section	controlled section and paragraph destination
claim text	atomic sentence or proposition
claim class	observed, represented, derived, determined, judgement or unresolved
support	evidence IDs, calculation IDs and exact citations
date and perimeter	applicable period, entities and scenario
permitted wording	qualification, attribution and uncertainty language
numeric check	canonical value, units and calculation version
permission check	allowed memo audience and purpose
reviewer	named role, disposition and date
invalidation	source, model, mandate or evidence change that requires refresh

APPENDIX D. MODEL RUN RECORD

Field	Entry
run ID	immutable identifier
task	classification, extraction, retrieval, comparison, explanation or draft
provider and model	exact provider, model and configuration
system and task instructions	versioned prompts and policy
tools	names, versions, permitted parameters and results
source population	permissioned source IDs, versions and exclusions
retrieval record	query, filters, ranking, returned context and citations
output	immutable candidate output
deterministic validation	schema, types, arithmetic, permissions and policy gates
evaluation	golden-set version, measures, thresholds and result
reviewer	identity, role, disposition, edits and date
downstream use	memo, question, analysis, exception or monitoring object
refresh	model, prompt, source, mandate or workflow trigger