

Underwriting Copilots for GCC Private Credit

A controlled architecture for jurisdiction evidence, escrow waterfalls, covenant monitoring and reviewable credit decisions

Prepared by Matchpoint Partners

Author attribution pending CK approval

August 2026

ABSTRACT

Background. GCC private-credit underwriting connects borrower and sponsor evidence, cash-flow normalisation, jurisdiction and security analysis, controlled accounts, covenants and portfolio monitoring. **Objective.** This paper develops an underwriting-copilot operating model for A3 private-credit, direct-lending and special-situations funds and A1 international institutional allocators. **Approach.** The analysis reviews 36 primary and authoritative sources available through 2 August 2026 and joins immutable sources, atomic evidence, an evidence graph, deterministic calculations, jurisdiction adapters, Claude-assisted cited work and named human approvals. **Findings.** A copilot can support extraction, reconciliation, comparison, explanation and drafting when calculations, permissions, exceptions and consequential decisions remain under approved controls. Enforcement is treated as evidence coverage and specialist routing rather than a model verdict. **Implications.** Escrow waterfalls and covenant tests should be versioned deterministic objects with exact lineage. The worked case, weights, thresholds, timings and economics 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. A3 private-credit funds and A1 international institutional allocators.

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

Architecture. Controlled sources feed an evidence graph, deterministic models, bounded Claude tasks, exceptions and authorised workflow.

Legal boundary. Jurisdiction adapters assemble evidence and questions; qualified counsel determines legal effect and enforceability.

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

JEL Classification: C88, G21, G23, G24, G28, G32, G33, K22, K25, O33

Keywords: GCC private credit, direct lending, underwriting copilot, Claude, covenant monitoring, escrow waterfall, enforcement evidence, security, collateral, evidence graph, agentic workflows

Research paper only. This document is not investment, legal, regulatory, compliance, accounting, audit, tax, valuation, privacy, cybersecurity or technology advice. Its worked case, score weights, thresholds, times and economics are unverified illustrative scenarios.

1. INTRODUCTION


Private-credit underwriting is an evidence problem before it is a modelling problem. A lender must identify the borrower and sponsor, reconstruct historical performance, normalise cash flows, test debt capacity, understand collateral and security, translate negotiated protections into monitorable covenants, and decide which unresolved matters must remain conditions precedent. In the Gulf Cooperation Council, that work also requires a clear jurisdiction and structure map. An onshore UAE operating company, a Dubai International Financial Centre vehicle, an Abu Dhabi Global Market holding company, a project escrow account and movable collateral can place different documents, registries, forums and specialist decisions into the same credit file.

This paper develops an underwriting-copilot operating model for A3 private-credit, direct-lending and special-situations funds and A1 international institutional allocators evaluating GCC private-market exposure. The design uses Claude as a bounded language and reasoning component inside a larger evidence system. It joins controlled source ingestion, document and table extraction, deterministic calculations, jurisdiction-specific rule packs, an evidence graph, exception workflows, human approvals and monitoring. The copilot prepares cited work and exposes uncertainty. Investment, legal, compliance, valuation, risk and disbursement authorities retain their existing decisions.

The central proposition is that underwriting speed should be measured at the point where a reviewable credit decision becomes possible. Fast document summarisation has little value when the output cannot be traced to the relevant source, period, calculation, structure or reviewer. A controlled copilot should therefore produce reusable decision objects: facts with exact citations; transformations with formulas and provenance; exceptions with owners; legal questions routed to qualified counsel; and release conditions tied to authorised evidence.

Research question	Operating answer
What should the copilot do?	Extract, reconcile, calculate, retrieve, compare, explain and draft within approved schemas and source permissions.
What should remain deterministic?	Arithmetic, date logic, covenant tests, approved waterfall logic, access rules, completeness tests and release gates.
How should enforcement risk be handled?	As a structured evidence-and-exception assessment by jurisdiction, forum, asset and structure; qualified counsel determines legal effect and enforceability.
How should escrow waterfalls be modelled?	From controlled account rules, transaction documents and verified cash events using a versioned calculation engine and approval workflow.
How should covenants be monitored?	Through event-driven agents that assemble evidence, run approved tests, classify exceptions and route review without changing rights or waiving breaches.
What proves value?	Observed, approved before-and-after operational evidence; this paper supplies only an unverified illustrative measurement design.

The framework draws on current UAE laws, Central Bank of the UAE standards, Dubai Land Department materials, current Dubai Financial Services Authority and Abu Dhabi Global Market sources, Basel Committee credit-risk principles, international financial-reporting standards, NIST AI guidance, W3C provenance standards, Anthropic product and evaluation materials, and the Financial Stability Board and International Monetary Fund discussion of private-credit vulnerabilities and data gaps [1-36]. Regulatory scope is stated at each use. Central Bank standards are directly applicable to licensed financial institutions within their scope and serve only as a benchmark for a private fund outside that perimeter. DFSA and ADGM rules apply only where the relevant entity, activity or fund is within their jurisdiction and authorisation.

Evidence-to-credit-decision chain**Figure 1. Evidence-to-credit-decision chain**

Source: Matchpoint underwriting-control framework; credit-risk context [1-3,23].

The paper makes no claim that a model can provide legal advice, determine enforceability, approve credit, waive a covenant, value collateral or release funds. The worked case, score weights, thresholds, timing assumptions and return-on-investment examples are [**Unverified illustrative scenarios**]. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence is supplied.

2. A3 AND A1 DECISION PERIMETER

2.1 A3 private-credit, direct-lending and special-situations funds

A3 is the operating user. Its investment team needs a defensible view of borrower quality, cash conversion, debt capacity, downside protection, documentation and monitoring. Portfolio operations and agency teams need the same underwriting objects to survive hand-off into the life of the loan. The fund may operate from outside the GCC, through a regional manager, or within a financial free-zone regime. Those organisational forms change the applicable permissions and governance. They do not change the requirement to connect every material conclusion to evidence and accountable authority.

The copilot should reduce repeated transcription and surface cross-document exceptions. It can compare management accounts with bank statements, map term-sheet covenants to draft agreements, trace a security package to required perfection evidence, and prepare a cited committee draft. It should never transform a plausible narrative into an accepted fact. Representations remain labelled representations until corroborated. Derived values remain labelled calculations with visible inputs and formula versions.

2.2 A1 international institutional allocators

A1 is an oversight and allocation user. An allocator assessing a GCC private-credit manager needs portfolio-level evidence about underwriting discipline, policy exceptions, covenant quality, jurisdiction concentration, valuation, arrears, amendments, watchlists and realised outcomes. It also needs a reliable distinction between the manager's reported data, independent evidence, system-derived metrics and specialist determinations.

An allocator-facing layer should therefore aggregate approved loan-level objects rather than generate a second ungoverned narrative. It may show the percentage of positions with complete security evidence, the distribution of jurisdiction lanes, average time to resolve material exceptions, override rates, covenant-data freshness and observed recovery outcomes. It should suppress borrower-confidential material outside the approved purpose and permission set. No allocator dashboard should reveal source content that the allocator is not authorised to receive.

2.3 Retained decision rights

Decision	Copilot contribution	Required authority
accept borrower and sponsor identity	extract and reconcile entities, identifiers and ownership claims	compliance and authorised onboarding owner
accept historical financials	reconcile source populations, periods and transformations	investment, finance or diligence owner
approve base and downside cases	calculate approved scenarios and display sensitivities	authorised investment committee
determine governing-law effect	retrieve relevant provisions and structure legal questions	qualified counsel in the relevant jurisdiction
determine enforceability	assemble evidence and unresolved conditions	qualified counsel; relevant court or authority ultimately applies law
approve collateral value	assemble valuation evidence and haircuts	authorised valuation and investment authority
set and waive covenants	draft and calculate approved terms	authorised lender parties under transaction documents
classify a breach or default	run tests and assemble event evidence	authorised agent, lender, counsel or committee as documents require
release a drawdown	verify recorded conditions and instructions	authorised facility, operations and payment controls
communicate to allocators	prepare permissioned, approved reporting	fund governance and information owner

A3 and A1 decision rights**Figure 2. A3 and A1 decision-rights map**

Source: Matchpoint ICP and decision-rights framework.

2.4 Three-state control language

Every material item should carry one of three states. **Supported** means the required evidence has been obtained, the source and period are identified, deterministic checks pass, and the designated reviewer has accepted the item for the stated purpose. **Conditioned** means the item can proceed only under a named condition, owner, due date and authority. **Unresolved** means the evidence is missing, conflicted, stale, out of permission or outside system competence. These states apply to data, calculations, security, covenants and legal questions. They avoid the false precision of a single confidence score.

3. GCC PRIVATE-CREDIT CONTEXT

3.1 Market significance and data limitations

The Financial Stability Board estimated the global private-credit market at approximately USD 1.5 trillion to USD 2 trillion in its May 2026 report and highlighted growing interconnections, leverage and data challenges [31]. The IMF's April 2024 Global Financial Stability Report discussed opacity, relatively fragile borrowers, valuation uncertainty and liquidity and leverage channels [33-34]. These are global observations. They do not establish the size, pricing, default rate or risk of the GCC segment. The Topic Tracker contains a coupon statement in the A3 legend; no approved current source for that statement was supplied, so this paper does not use it as a fact.

For underwriting, the practical implication is a need for transaction-level evidence. A broad market view cannot substitute for borrower cash flow, sponsor support, legal structure, collateral, covenant headroom and disbursement control. Portfolio aggregation should preserve those underlying definitions so that a concentration or exception ratio is not assembled from incompatible fields.

3.2 Regulatory lanes

The UAE has several relevant legal and regulatory lanes. UAE federal legislation includes commercial transactions, electronic transactions, evidence, personal-data protection, bankruptcy, commercial companies, civil transactions and movable-security provisions [4-11]. Dubai real-estate development escrow accounts are governed by Dubai legislation and implemented through the Dubai Land Department [14-16]. DIFC financial services fall within the DFSA framework; ADGM financial services and insolvency operate within the ADGM framework [17-22]. Other GCC countries have their own laws, registries, regulators and courts. A system may reuse a common data model across these lanes, but each jurisdiction adapter must be owned, reviewed, dated and approved by appropriate specialists.

The DFSA's current credit-fund page describes rules for credit funds, including fund form and portfolio requirements within the DIFC regime [17]. Those rules should inform a manager or fund only when the relevant vehicle and activity are in scope. An ADGM consultation on private-credit funds and a July 2026 DFSA consultation are proposals or consultation materials, not final operative rules for purposes of this paper [21-22]. The copilot must store status and effective date so that a proposal is never silently applied as law.

3.3 Credit-risk standards as an operating benchmark

The CBUAE Credit Risk Management Regulation and Standards require licensed financial institutions within scope to maintain sound underwriting, administration, measurement, monitoring, documentation, collateral and exception practices [1-2]. The CBUAE capital-adequacy standards also address legal certainty for credit-risk mitigation [3]. Basel's 2025 Principles for the Management of Credit Risk organise credit risk around an appropriate environment, sound granting processes, administration and monitoring, and adequate controls [23]. A non-bank private-credit fund may use these materials as an authoritative operating benchmark while clearly recording that benchmark use does not create direct regulatory applicability.

3.4 Unit of analysis

The underwriting system should distinguish at least five connected units:

1. the obligor and sponsor group;
2. the facility and each tranche or instrument;
3. each collateral asset and security interest;
4. each account, waterfall and payment event; and
5. each covenant, reporting duty, consent and exception.

This separation matters because risk and rights can diverge. A consolidated borrower group may generate cash in one entity, pledge an asset in another, route collections through a controlled account, and incur covenants at facility and obligor levels. A model that flattens these relationships can overstate available cash, double-count collateral or apply a covenant to the wrong perimeter.

4. UNDERWRITING EVIDENCE MODEL

4.1 Source hierarchy

Source class	Examples	Permitted use	Minimum control
authoritative public source	legislation, regulator rulebook, official registry record	legal and regulatory context; identity or filing evidence within scope	jurisdiction, status, effective date and retrieval date
executed transaction document	facility, intercreditor, security, escrow and account-control agreements	rights, duties, definitions, tests and conditions	complete signed version, amendment chain and qualified interpretation
controlled financial record	audited accounts, signed management accounts, general ledger and bank record	historical and current financial evidence	entity, period, currency, basis and reconciliation
independent specialist record	valuation, engineering, audit, legal opinion and administrator confirmation	specialist-controlled proposition	issuer, scope, date, assumptions and reliance limits
borrower or sponsor representation	data-room schedule, certificate, forecast and questionnaire	represented fact and scenario input	representative, date, warranty status and corroboration plan
lender or agent record	approval, waiver, notice, calculation and payment instruction	internal decision and transaction administration	authority, version, effective period and recipient
market or third-party dataset	prices, rates, adverse information and sector data	contextual analysis	provider, licence, timestamp, method and limitations
analyst or system work product	model, memo, score and exception	derived analysis	cited inputs, formula or configuration, reviewer and version

4.2 Atomic evidence objects

A credit memo paragraph is too large to serve as the primary data unit. Each material proposition should be captured as an atomic evidence object. The object identifies the subject, predicate and value; the entity and facility perimeter; the source and exact location; the period; currency and units; source authority; whether it is observed, represented, derived or determined; transformation history; permission; conflict state; reviewer; and refresh trigger.

For example, “EBITDA was AED 42 million” is incomplete. The object should state which entity or consolidation perimeter, which twelve-month period, which accounting basis, whether the number came from audited or management accounts, which normalisations were applied, who accepted those normalisations and which downstream leverage tests used the value. The model may explain the number; the calculation record establishes it.

4.3 Evidence graph

The evidence graph connects entities, facilities, assets, accounts, documents, claims, calculations, covenants, events and decisions. W3C PROV-O provides a standard vocabulary for relating entities, activities and agents [30]. A practical credit graph should also store bi-temporal information: when a claim was true in the transaction world and when the system learned or changed it. This allows a committee to reproduce what evidence was available at approval and to distinguish a late-arriving correction from the original decision record.

Underwriting evidence graph and copilot architecture




Figure 3. Underwriting evidence graph and copilot architecture

Source: Matchpoint architecture; Claude and provenance context [26-30].

4.4 Completeness matrix

Before analysis begins, the lender should define a requirements matrix by deal type, structure and stage. Rows represent required evidence; columns record applicability, expected provider, due date, receipt, version, review owner, exceptions and release dependency. The matrix prevents a strong narrative from hiding a missing legal opinion, unsigned security instrument or incomplete bank population.

Typical populations include corporate and ownership evidence, financial statements, bank and ledger data, budgets and forecasts, debt schedules, tax and regulatory records, material contracts, insurance, collateral and valuation materials, existing security and priority evidence, litigation, compliance, term sheet, facility and security documents, escrow or account-control documents, covenant definitions, conditions precedent and monitoring reports.

4.5 Negative and missing evidence

The absence of a document is not evidence that the underlying fact is false. It is an unresolved evidence condition. “No litigation found in the supplied materials” differs from “the borrower has no litigation.” The first is a bounded search result; the second requires a source and authority capable of supporting it. The copilot should state the population searched, the date, the method and the remaining sources or confirmations required.

5. COPILOT ARCHITECTURE

5.1 Architecture layers

Layer	Function	Control
source systems	data room, ledger, bank, CRM, registry, legal and servicing records	owner, permission, version, retention and immutable original
ingestion	receipt, malware scan, OCR, table extraction and classification	file hash, completeness, quality and chain of custody
evidence store	source text, tables, images, coordinates and metadata	exact citation, period, access and deletion policy
evidence graph	entities, claims, facilities, assets, accounts, covenants and decisions	typed schema, bi-temporal provenance and reversible corrections
calculation engine	financial normalisation, ratios, sensitivities, waterfall and covenant tests	versioned formulas, units, input lineage and independent checks
retrieval	lexical, structured and graph retrieval	permission before retrieval, authority ranking and bounded context
Claude	classification, extraction, comparison, explanation and drafting	structured output, citations, abstention, evaluation and prompt boundary
rules and workflow	requirements, exceptions, approvals, notices and release gates	named owners, segregation of duties and immutable decision log
monitoring	quality, drift, access, exceptions, cost and incidents	threshold, alert, fallback, rollback and review cadence

Claude's citation feature and tool-use interfaces can support cited responses and calls to approved deterministic services [26-27]. Those product capabilities do not prove that a particular deployment is accurate, secure or suitable. Each configuration, tool, schema, prompt and retrieval method requires task-specific evaluation. Anthropic's work on evaluating AI systems and trustworthy agents supports evaluation and control thinking; it does not replace the lender's governance or independent validation [28-29].

5.2 Task decomposition

A broad instruction such as “underwrite this borrower” is unsuitable. The workflow should decompose underwriting into bounded tasks with explicit inputs and outputs. Examples include classifying a document, extracting a defined table, reconciling entity names, linking a figure to an audited note, calculating a ratio using an approved formula, comparing a covenant definition across drafts, retrieving evidence for an exception, and drafting a paragraph from accepted evidence objects.

Each task should declare materiality, allowed sources, required fields, validation rules, abstention conditions and human authority. The model should return null or an explicit exception when a required field is absent. It should not fill a missing month, infer a guarantee or convert an unexecuted draft into an obligation.

5.3 Retrieval contract

Retrieval should filter by permission, entity, facility, document status, effective date, source hierarchy and task. A covenant calculation should retrieve the executed definition and current amendments before drawing on a committee summary. A legal question should retrieve the relevant instrument and official sources, then route the question to counsel. A portfolio report should retrieve approved aggregates rather than raw borrower documents.

The retrieval record should preserve the query, filters, returned sources, ranking, model context and exclusions. Citation presence is only a first control. Citation support must be tested: the cited passage must support the complete atomic proposition, apply to the right entity and period, and have sufficient authority for the claim.

5.4 Deterministic tool boundary

The language model should call controlled tools for arithmetic, currencies, dates, covenant definitions, financial models, registry lookups and document comparison. The tool layer should validate units, reject incompatible periods, log formulas and produce machine-readable exceptions. Natural-language explanations can then be generated from the approved outputs. This separation preserves reproducibility and limits the risk that persuasive language obscures a calculation error.

5.5 Configuration record

Every production run should identify the model and provider, configuration date, system instructions, task prompt, schema, tools, retrieval policy, source population, temperature or equivalent settings, output, validation results, reviewer

and downstream use. Model changes should trigger regression tests against a versioned golden set. The system should be able to reconstruct which configuration produced a committee paragraph or monitoring alert.

6. INTAKE AND TERM-SHEET NORMALISATION

6.1 Intake perimeter

The intake record should define the proposed borrower group, sponsor, facility type, amount, currency, use of proceeds, maturity, pricing components, repayment, security, guarantees, accounts, governing law, forum, conditions, covenants, information undertakings and timetable. The record should distinguish borrower-provided information from lender-authored terms and adviser summaries.

An entity map comes first. It links legal names, registration identifiers, jurisdictions, ownership and control, operating roles, asset ownership, contracting roles and proposed obligors. Entity resolution should use durable identifiers where available. Similar names are candidate matches, not proof of identity. Material conflicts become exceptions that block aggregation.

6.2 Term extraction

The copilot may extract term-sheet fields into a controlled schema. It should retain exact source locations, qualifiers and nulls. It should detect linked definitions, options, step-ups, baskets, cure rights and documentary dependencies. A term expressed as “subject to agreed documentation” remains provisional. A term in a lender presentation is not automatically part of the executed facility.

Term object	Required fields
commitment	amount, currency, lender, availability and reduction events
pricing	base rate, floor, margin, fees, step-up, day-count and payment dates
repayment	amortisation, bullet, cash sweep, prepayment and cancellation
use of proceeds	permitted uses, exclusions, evidence and draw control
security	grantor, asset, jurisdiction, priority, perfection and release
guarantee	guarantor, scope, limit, governing law and conditions
covenant	definition, threshold, test date, perimeter, cure and consequence
reporting	item, period, due date, form, signatory and delivery channel
condition	evidence, responsible party, satisfaction authority and waiver route
account control	bank, account, permitted flows, waterfall and instruction rights

6.3 Draft comparison

Document comparison should operate at clause and defined-term levels. The system should identify additions, deletions, changed thresholds, altered definitions, cross-reference changes and silent dependencies. It should avoid declaring that two clauses have the same legal effect. Counsel and authorised deal teams determine materiality and legal effect. The output should give both source passages, the structured difference and the affected calculations or conditions.

6.4 Requirements generation

The normalised structure can generate a deal-specific evidence matrix. A real-estate development loan with escrowed sales receipts will require a different account and project evidence population from an acquisition facility secured over shares and receivables. The system should use approved templates as a starting point, then require reviewers to confirm applicability and additions. Template coverage is never evidence of transaction completeness.

7. BORROWER, SPONSOR AND CASH-FLOW ANALYSIS

7.1 Financial population

Financial ingestion should preserve entity, consolidation perimeter, accounting basis, currency, unit, period and source status. The system should reconcile audited financial statements, signed management accounts, general-ledger extracts, bank records, tax filings and operating data where supplied and permitted. Differences should be explained or left unresolved; a model should not silently choose the number that fits the requested narrative.

IFRS 9 addresses classification, measurement and expected credit losses for financial instruments [24]. Its direct applicability depends on the reporting entity and accounting framework. A lender may use IFRS 9 concepts within its approved provisioning and risk process while retaining its own underwriting and monitoring evidence. The copilot should not assign an accounting classification or impairment stage without the authorised accounting policy and reviewer.

7.2 Normalisation bridge

Each normalisation should be a separate adjustment object with source, rationale, period, amount, sign, recurrence status, treatment by case, reviewer and date. The bridge from reported earnings to underwriting cash flow should be reproducible. Management's forecast adjustment remains a representation. A lender adjustment becomes a decision only when approved under policy.

Adjustment test	Question	Control
entity	Does the item belong to the obligor or support provider?	map to legal entity and consolidation perimeter
period	Is it inside the tested period?	retain transaction and accounting dates
cash	Does it affect available cash?	reconcile to cash flow and working capital
recurrence	Is there evidence that it is exceptional?	require pattern and documented basis
control	Can the borrower influence the item?	distinguish committed action from intention
downside	How does the item behave under stress?	define scenario-specific treatment
covenant	Does the executed definition permit it?	calculate separately from underwriting view

7.3 Debt-capacity engine

The engine should maintain distinct base, downside and severe-but-plausible cases. Each case records commercial assumptions, source, owner and approval. Calculations should include debt service, interest, fees, mandatory amortisation, cash taxes, working capital, maintenance capital expenditure, permitted distributions, trapped cash and liquidity. The system should expose the difference between accounting earnings, covenant EBITDA and cash available for debt service.

Debt capacity should not be reduced to one ratio. Relevant outputs may include gross and net leverage, fixed-charge coverage, debt-service coverage, interest coverage, minimum liquidity, loan-to-value, borrowing-base coverage, cash-conversion and maturity concentration. Definitions vary by transaction. Every displayed ratio should link to its approved formula, exact inputs and test perimeter.

7.4 Sponsor analysis

Sponsor support requires evidence of capacity, willingness, legal obligation and availability. Reputation or past support is contextual evidence, not a guarantee. A binding guarantee, equity commitment or keepwell arrangement requires document-specific legal analysis. The copilot can assemble ownership, fund, liquidity, track-record and support evidence; authorised investment and legal reviewers determine credit value.

7.5 Forecast challenge

The copilot can compare forecast assumptions with historical ranges, contracts, capacity, working-capital cycles and sector variables. It should generate questions and sensitivities rather than assert a forecast. A good output states which assumption drives covenant headroom, when liquidity turns negative, which operational evidence supports the assumption, and what management information will monitor it after closing.

8. ENFORCEMENT-RISK EVIDENCE MODEL BY JURISDICTION AND STRUCTURE

8.1 Purpose and boundary

Enforcement risk cannot be captured as a model verdict. Outcomes depend on governing law, forum, obligor and asset location, security type, creation and perfection, priority, evidence, insolvency status, procedural events and the applicable authority. The copilot should therefore produce an evidence-coverage and exception model. Qualified counsel in the relevant jurisdiction determines legal effect, enforceability and required remedial action.

The model uses a jurisdiction lane and a structure lane. For the UAE, lanes may include federal onshore law and the relevant local emirate judicial and registry context, DIFC, and ADGM. A loan may touch several lanes at once. Other GCC states require separately approved country adapters. The shared schema remains constant while official sources, registries, forms, terminology, statuses and counsel ownership vary.

8.2 Evidence dimensions

Dimension	Evidence question	Example exception	Authority
parties	Are obligors, grantors and secured parties correctly identified and authorised?	incompatible identifier or missing authority	compliance and counsel
governing law	Which law governs each obligation and security instrument?	summary conflicts with executed clause	counsel
forum	Which court or arbitration route is agreed and available?	inconsistent jurisdiction clauses	counsel
asset situs	Where is each asset located, registered or controlled?	asset location unverified	counsel and asset specialist
creation	Does the instrument contain the required grant and description?	draft or incomplete schedule	counsel
perfection	Which filing, possession, control, notice or registration is required?	filing receipt absent or stale search	counsel and registry owner
priority	What prior interests, statutory claims or intercreditor terms affect ranking?	unresolved competing filing	counsel
evidence	Can execution, signature, notice and record integrity be established?	incomplete signature or source chain	counsel and evidence owner
insolvency	Which stay, avoidance, restructuring or distribution rules may apply?	cross-border group path unresolved	insolvency counsel
operations	Can accounts, assets and information be controlled in practice?	contractual right lacks operational setup	agent, operations and counsel

UAE electronic-transactions and evidence legislation provides a legal framework for electronic records, documents and signatures [5-6,9]. The current UAE Bankruptcy Law and its executive regulation govern federal bankruptcy processes within their scope [7-8]. The UAE Commercial Transactions, Commercial Companies and Civil Transactions laws provide additional context [4,9-10]. These sources should be retrieved by effective date and used by qualified specialists. The system should not translate their existence into an enforceability conclusion.

Jurisdiction-and-structure enforcement evidence matrix**Figure 4. Jurisdiction-and-structure enforcement evidence matrix**

Source: Matchpoint legal-readiness framework; UAE, DIFC and ADGM context [3-22].

8.3 Scoring as triage

An illustrative score can help prioritise incomplete evidence. It should measure coverage and unresolved consequence, not the probability that a court will enforce. Hard-stop conditions remain separate. A sample design allocates points across party and authority evidence, instrument completeness, perfection, priority, forum and governing-law coherence, insolvency analysis and operational control. A high coverage score cannot override a missing registration or adverse counsel conclusion.

Each score component should display its evidence, rule version, reviewer and expiry. The label should be “enforcement evidence coverage” or “legal-readiness triage,” followed by an explicit statement that it is not a legal opinion. Material score changes should be explainable at the evidence-object level.

8.4 Counsel workflow

The copilot should convert exceptions into bounded legal questions. For example: identify the grantor, asset, governing law, proposed perfection step, available filing evidence, competing interests and requested conclusion. Counsel's response should be stored as a determination object with scope, assumptions, date, reliance limitation and required actions. The system then links the determination to conditions precedent and post-closing obligations.

8.5 Currency and change control

Legal adapters need named owners and review dates. Official-law changes, court developments, registry changes, instrument amendments and asset movements can trigger refresh. Consultation documents should be labelled proposal or consultation. Archived rules must remain available to reproduce historical decisions. A current page should never be assumed to describe the rules effective at an earlier closing date.

9. COLLATERAL, SECURITY AND PERFECTION

9.1 Collateral register

The collateral register should identify each asset, legal and beneficial owner, jurisdiction, location, description, valuation, currency, eligible amount, haircut, existing interest, proposed security, perfection method, evidence, reviewer and monitoring trigger. Links should connect the asset to facility, obligor, security instrument and relevant account or insurance.

The UAE movable-assets executive regulation and the Emirates Movable Collateral Registry provide an official framework and registry service for movable collateral within scope [12-13]. Registry evidence should record the exact search or filing, names and identifiers used, date and result. A search under one spelling or identifier does not prove absence under all relevant identities. Counsel determines the required search and filing population.

9.2 Valuation and eligibility

Valuation is a specialist process. The copilot may reconcile valuation dates, methods, assumptions, currencies, ownership and asset descriptions; it may calculate policy haircuts using approved rules. It should not originate a collateral value from unverified narrative. Valuation age and trigger events should feed an exception queue.

Borrowing-base eligibility should be deterministic. Each asset either satisfies the approved definition, is ineligible, or remains conditioned on evidence. Concentration limits, reserves and advance rates must show formula and source. The system should prevent double-counting across facilities or collateral pools.

9.3 Security checklist

Control stage	Required record
scope	asset, grantor, secured obligations and security type
authority	constitutional, board and signatory evidence
execution	complete instrument, signature and date evidence
perfection	filing, notice, possession, control or registration requirement
priority	searches, prior interests, releases and intercreditor terms
opinion	counsel scope, assumptions, qualifications and conclusion
custody	location and controller of originals or control credentials
release	permitted release event, authority and evidence
monitoring	expiry, renewal, asset movement, valuation and insurance triggers

9.4 Conditions precedent and subsequent

Conditions should be machine-readable objects. Each has a description, source clause, evidence requirement, obligor, owner, due date, satisfaction authority, waiver authority, dependencies and status. The model can assemble and compare evidence; only the designated authority marks satisfaction or waiver. Conditions subsequent need the same discipline because an open post-closing item can alter priority or recovery.

10. ESCROW-WATERFALL MODELLING

10.1 Legal and operational context

Dubai Law No. 8 of 2007 establishes a framework for escrow accounts for real-estate development in Dubai [14]. Dubai Land Department materials describe services and frequently asked questions concerning project escrow and mortgage payments into escrow accounts [15-16]. Applicability and permitted flows depend on the project, account, transaction documents and current official requirements. The copilot should retrieve and label those sources, then route the legal and operational design to appropriate DLD, bank, counsel and lender authorities.

An escrow waterfall model should reproduce the agreed order of cash application. It should not infer that cash is freely available merely because it appears in an account. Each cash event needs account identity, value date, currency, payer or beneficiary, transaction reference, permitted purpose, source evidence, reconciliation status and approval state.

10.2 Waterfall objects

Object	Minimum fields
account	bank, account identifier, currency, legal owner, control and purpose
cash event	amount, value date, source, payer, beneficiary and evidence
rule	priority, trigger, calculation, cap, reserve and destination
reserve	target, balance, permitted use, replenishment and release
instruction	authorised party, date, amount, beneficiary and verification
exception	unmatched receipt, shortfall, prohibited purpose, timing or data gap
approval	authority, scope, conditions, date and evidence
output	applied amount, residual cash, next priority and unresolved items

10.3 Deterministic calculation

The waterfall engine should ingest verified cash events and apply a versioned rule set. Rules can allocate taxes and statutory items where applicable, project costs, operating costs, reserve replenishment, fees, interest, scheduled principal, mandatory prepayment and permitted distributions in the documented order. The exact order is transaction-specific. Each output should link back to the executed clause or approved operating rule.

Currency conversion, day count, caps, thresholds and pro rata allocations should be deterministic. The engine should reject missing exchange-rate sources, incompatible value dates, duplicate transactions and unauthorised beneficiaries. The language model may explain the result and identify the clauses used.

Controlled escrow-waterfall calculation




Figure 5. Controlled escrow-waterfall calculation

Source: Matchpoint calculation framework; Dubai escrow context [14-16].

10.4 Reconciliation and release

At least three populations may require reconciliation: bank transactions, project or borrower records, and agent or lender calculations. Differences should become named exceptions. A release pack should show opening balance, verified receipts, applied waterfall, reserved amounts, proposed payments, beneficiary verification, residual cash, exceptions and required approvals. A separate authorised payment workflow executes the release.

10.5 Scenario analysis

Scenario modelling can test delayed collections, lower sales, cost overruns, reserve leakage, interest-rate changes and maturity pressure. These are forecasts, not observed outcomes. The model should display the assumption source and owner, identify the first breached threshold, and state which facility or operational right may become relevant subject to document and legal review.

11. COVENANT ARCHITECTURE

11.1 Covenant as executable specification

A covenant object should contain the full definition, threshold, direction, testing perimeter, frequency, test date, look-back period, currency, formula, permitted adjustments, baskets, cure mechanics, grace period, reporting evidence, source clauses, amendment history and authorised interpretation. The executed document controls. A spreadsheet or prior monitoring note is a derived implementation that must reconcile to it.

The system should distinguish maintenance covenants, incurrence tests, information undertakings, affirmative and negative undertakings, conditions, events of default and operational triggers. They may share data but have different legal consequences. A calculated variance does not by itself establish a breach or default.

11.2 Definition graph

Defined terms often depend on other terms, schedules and exceptions. The copilot should construct a definition graph and flag circular references, missing definitions, draft changes and dependencies. Counsel and authorised deal teams validate interpretation. The approved formula should encode the accepted operational reading and link to the determination.

11.3 Data mapping

Covenant component	Data source	Validation
tested period	reporting calendar and delivered certificate	dates and completeness
consolidation perimeter	entity map and definition	approved inclusions and exclusions
earnings input	financial statements and normalisation bridge	source status and permitted add-backs
debt input	facility, bank and debt schedule	instrument and currency reconciliation
cash input	controlled account and bank evidence	availability and restriction test
valuation input	approved specialist record	date, scope, currency and haircut
threshold	executed covenant and amendments	effective-date control
result	deterministic calculation	independent check and version
consequence	executed documents and determination	authorised legal and facility review

11.4 Amendment and waiver control

Amendments, consents and waivers must update the definition and rule versions for their effective period. The original remains available for historical reproduction. A waiver may be limited to a period or event; it should not silently change future calculations. The agent should route an expired or inapplicable waiver as an exception.

11.5 Covenant certificate workflow

The borrower certificate, underlying financial evidence and lender calculation should be separate records. The system compares them, explains variances and records the reviewer. Signature, authority and delivery evidence should be captured. An accepted certificate does not remove the need to preserve the underlying evidence and lender decision where policy requires it.

12. COVENANT-MONITORING AGENTS

12.1 Agent boundary

A covenant-monitoring agent is an event-driven workflow that retrieves permitted evidence, invokes approved tools, assembles a calculation and routes exceptions. It has no authority to amend a facility, accept a waiver, classify a default, contact a borrower or release funds unless a separately authorised workflow explicitly permits that action. The system should start with read and draft permissions.

12.2 Trigger model

Triggers may include a reporting due date, document receipt, bank event, account shortfall, valuation expiry, insurance expiry, ownership change, missed payment, covenant headroom threshold, amendment, adverse-information review or manual request. Every trigger should create a case with scope, source, time and owner. Duplicate and superseded triggers should be resolved deterministically.

12.3 Agent workflow

1. identify the facility, obligor, test and effective rule version;
2. check source permissions and required evidence population;
3. retrieve accepted and newly supplied evidence;
4. validate entity, period, currency, units and completeness;
5. call deterministic calculations and reconciliation tools;
6. compare borrower, agent and lender results;
7. classify data or calculation exceptions under approved rules;
8. draft a cited review pack and proposed next actions;
9. route to the authorised reviewer; and
10. record the reviewer decision, communication and refresh dependencies.

Covenant-monitoring agent workflow




Figure 6. Covenant-monitoring agent workflow

Source: Matchpoint monitoring architecture; credit-risk and agent context [1-3,23,27-29].

12.4 Exception taxonomy

Exception	Example	Default route
missing evidence	management accounts or certificate absent	facility operations and deal team
stale evidence	valuation or insurance beyond approved date	collateral owner and deal team
identity mismatch	entity or account differs from approved map	compliance and operations
period mismatch	certificate uses a different look-back period	finance and facility agent

Exception	Example	Default route
definition mismatch	add-back not permitted by approved rule	investment, finance and counsel
calculation variance	borrower and lender results differ	finance and deal team
threshold event	calculated headroom crosses alert level	portfolio risk and deal team
potential document event	observed fact may engage a right or default	counsel and authorised facility parties
permission conflict	source is outside recipient purpose	information owner and privacy
model uncertainty	retrieval or extraction cannot support an answer	human review; no automated conclusion

12.5 Communications

The agent may draft a borrower query, reservation-of-rights note, committee update or allocator summary only from approved evidence and templates. The draft must identify its status and required approvers. Sending remains a separately authorised action. Legal communications require counsel review where appropriate. The system should record the evidence version used because a later correction may require a revised communication.

12.6 Audit and supervision

Agent actions, tool calls, retrieved sources, calculations, proposed classifications, approvals and communications should be logged. Supervisors should be able to inspect sampling results, override patterns, delayed cases, false alerts, missed events and access incidents. High automation volume is not a quality measure. Relevant measures concern correct evidence, timely review and preserved decision rights.

13. EARLY WARNING AND EXCEPTION MANAGEMENT

13.1 Early-warning design

Early-warning indicators should be defined as observed values or approved transformations. Examples include declining headroom, liquidity reduction, delayed receivables, collection variance, account leakage, repeated reporting delay, customer concentration, cost overrun, valuation decline and requests for amendment. An indicator should state its source, direction, threshold, look-back period, expected lag and owner.

The system should avoid implying causation. A delayed report may reflect operating stress, process weakness or a harmless administrative delay. It is a trigger for review. The review record captures the explanation and evidence.

13.2 Watchlist case

A watchlist case connects indicators, exceptions, borrower interactions, specialist advice, scenario results, decisions and actions. It should show current liquidity, debt service, covenant headroom, collateral and upcoming events using approved data. Each action has an owner, due date and evidence requirement. Closed items remain in history.

13.3 Portfolio aggregation

For A1 reporting, loan-level objects can be aggregated into jurisdiction, sector, sponsor, instrument, maturity, covenant and exception views. Aggregation requires common definitions and an as-of date. The system should report coverage alongside the metric. For example, a covenant-headroom distribution should state the proportion of portfolio value with current, accepted calculations.

13.4 Escalation

Escalation rules should combine severity, materiality, time and authority. A missing routine report differs from an unverified payment instruction or absent perfection evidence. Some events need immediate hard-stop routing regardless of a score. The rules should be approved, versioned and tested against historical cases.

14. WORKED UNDERWRITING SCENARIO

14.1 Scenario status

The following case is an **[Unverified illustrative scenario]**. It does not describe a Matchpoint client, observed transaction or legal conclusion. Names, values, timings, probabilities and outcomes are invented solely to demonstrate the framework. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0.

14.2 Proposed facility

Gulf Project Holdings Ltd seeks a four-year AED 180 million senior secured facility for a UAE operating and development group. The proposed security package includes shares in an onshore obligor, receivables, specified movable assets and controlled collection accounts. Certain project receipts are represented as subject to a Dubai project escrow arrangement. A sponsor support undertaking is proposed. The facility includes leverage, debt-service coverage, minimum-liquidity, account and information covenants.

The supplied data room contains audited accounts for two years, current management accounts, bank statements, a financial model, project contracts, a draft term sheet, corporate documents, a valuation, a security schedule and draft account instructions. The copilot labels each source by status and builds the requirements matrix.

14.3 Evidence exceptions

Entity resolution identifies three spellings for one operating company and an incompatible registration identifier in a security schedule. The system blocks the proposed asset link pending verification. Bank reconciliation finds a material collection account outside the supplied account map. The valuation describes assets using identifiers that do not match the security schedule. A sponsor-liquidity statement is supported only by a management representation.

Financial normalisation identifies an AED 12 million forecast add-back. The approved covenant definition does not yet permit it, and the supporting contract is unsigned. The item remains excluded from the covenant calculation and appears separately in the investment sensitivity. The downside case shows debt-service coverage below the illustrative committee threshold in two quarters when sales receipts are delayed by sixty days.

14.4 Jurisdiction and security review


The structure map separates onshore UAE entities and assets, the project escrow, and any free-zone holding or finance entity. The legal-readiness assessment shows missing evidence for one movable-security filing, incomplete authority for one grantor and an unresolved account-control mechanism. The copilot generates bounded counsel questions and links responses to conditions precedent. It does not assign an enforceability probability.

14.5 Waterfall model

The illustrative waterfall uses verified opening balances and labelled forecast receipts. It allocates permitted project costs, reserve replenishment, interest and scheduled principal under the assumed rule version. A proposed distribution is blocked in the model because the minimum reserve and downside headroom are not met. An authorised reviewer determines the transaction response.

14.6 Committee pack

The final pack separates supported facts, representations, approved calculations and unresolved conditions. It includes the entity and facility map, source coverage, normalisation bridge, base and downside cases, security evidence, account waterfall, covenant definitions, legal questions and conditions precedent. The recommendation is a committee decision, not a model output.

Worked underwriting scenario**Figure 7. Worked underwriting scenario from intake to conditioned decision**

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

14.7 Monitoring hand-off

At closing, accepted underwriting objects are promoted into monitoring. Covenant formulas, account rules, reporting requirements, security refresh dates and conditions subsequent retain their versions and owners. Unresolved or waived items preserve their scope and expiry. The portfolio team receives the same evidence lineage used by the committee.

15. GOLDEN-SET EVALUATION

15.1 Evaluation population

The golden set should represent document and transaction diversity: audited and management accounts, bank statements, term sheets, facility and security instruments, account agreements, valuations, certificates, amendments, waivers, registry evidence and monitoring reports. It should include poor scans, bilingual or multilingual materials where within scope, inconsistent tables, repeated names, amended definitions, missing schedules and adversarial instructions embedded in documents.

Every expected answer should have an authoritative source and reviewer. Where specialists disagree, the set should record the acceptable boundary and abstention requirement rather than force artificial consensus.

15.2 Component metrics

Component	Primary measure	Material failure
classification	precision and recall by document class	critical document missed or mis-stated
extraction	field accuracy with source-span accuracy	party, amount, date, currency or threshold error
table recovery	cell and structure accuracy	detached header or shifted period
entity resolution	pairwise precision and recall	false merge across legal entities
retrieval	authoritative evidence recall at bounded context	executed clause omitted
citation	atomic-claim support and exact location	unsupported material proposition
calculation	agreement with approved deterministic result	ratio, waterfall or currency error
comparison	reviewed change detection	amendment or threshold change missed
abstention	correct abstention under missing or conflicted evidence	confident unsupported answer
workflow	correct exception, owner and approval path	decision-right bypass

15.3 End-to-end tests

End-to-end evaluation should ask whether a reviewer can reach the correct controlled state using the pack. Tests include a supported underwriting conclusion, a conditioned conclusion, a rejected or deferred case, a security exception, a covenant variance, an escrow shortfall, an unauthorised source and a model-injection attempt. The expected result includes both content and routing.

15.4 Severity weighting

Errors should be weighted by consequence. A punctuation difference has little importance. A false entity merge, missing negative covenant, wrong currency, omitted filing condition or unsupported payment instruction may be critical. Release criteria should combine aggregate measures with zero-tolerance categories for high-consequence failures.

15.5 Production monitoring

Monitoring should record task volume, exceptions, reviewer overrides, citation failures, calculation failures, model and prompt version, latency, cost, access denial, incident and rollback. Drift should be assessed by document type, language, transaction structure and jurisdiction lane. A stable average can conceal a severe failure in a small but material class.

16. SECURITY, PRIVACY AND MODEL GOVERNANCE

16.1 Data protection

The UAE Personal Data Protection Law establishes requirements relevant to processing personal data within its scope [11]. Transaction materials can also include confidential, privileged, commercially sensitive and regulated information. The system should classify data, record purpose and permission, minimise retrieved content, encrypt data, segregate environments, manage retention and deletion, and log access.

The CBUAE Outsourcing Regulation and enabling-technologies guidance apply to licensed financial institutions within their respective scope and provide authoritative control context for outsourcing and technology adoption [25,36]. A private fund outside that scope may use them as a benchmark and should label that status. Supplier contracts, data location, subcontractors, incident duties, audit rights, exit and deletion require owner approval.

16.2 Access before retrieval

Permission must be enforced before content enters model context. Post-generation redaction is insufficient for unauthorised retrieval. Source objects should inherit transaction, entity, purpose, recipient and confidentiality controls. Portfolio reporting should draw from approved aggregates. Privileged materials require a separately approved route.

16.3 Prompt injection and untrusted content

Documents, websites and messages are untrusted content. An instruction embedded in a borrower document should be treated as text evidence, not as a system command. Tools should be allow-listed, typed and scoped. High-impact actions should require explicit workflow authority. The system should validate every tool response and keep a complete call record.

16.4 AI risk management

The NIST AI Risk Management Framework Generative AI Profile and AI Resource Center provide voluntary risk-management resources [32,35]. A lender can map governance, measurement and management controls to its own risk framework. The mapping should identify owners, evidence, thresholds and residual risk rather than create a compliance badge.

16.5 Human factors

Reviewers need enough time, source access and authority to challenge the output. Interfaces should expose conflicts and missing evidence prominently. Automation bias should be tested through seeded errors and blind review. Overrides should require a reason and feed evaluation. A model-generated committee draft should be clearly labelled until authorised.

16.6 Incident and fallback

The operating plan should cover incorrect output, data leakage, access failure, supplier outage, source corruption, calculation defect and unauthorised action. Fallback procedures should preserve manual underwriting and monitoring. Incidents should identify affected outputs, transactions, recipients and decisions, then trigger correction and notification under approved policy.

17. NINETY-DAY IMPLEMENTATION ROADMAP

17.1 Days 0 to 15: scope and authority

Select one bounded facility type and one jurisdiction lane. Name executive, investment, credit, legal, compliance, information-security, privacy, operations, model-risk and data owners. Define retained decisions, prohibited actions, source permissions, materiality and success measures. Freeze the first requirements, evidence and covenant schemas.

17.2 Days 16 to 30: evidence foundation

Implement immutable intake, file hashing, document status, entity and facility mapping, exact citations and the requirements matrix. Build deterministic extraction validation for dates, currencies, periods and identifiers. Establish the exception record and reviewer workflow.

17.3 Days 31 to 45: calculations and legal-readiness lane

Encode an approved financial model, normalisation bridge and a limited covenant set. Build one jurisdiction-and-structure evidence adapter with counsel ownership. Create the security and condition registers. Validate output against historical reviewed cases.

17.4 Days 46 to 60: Claude in shadow mode

Enable bounded extraction, comparison and cited drafting. The existing underwriting process remains authoritative. Capture false positives, misses, unsupported claims, abstentions, access failures, reviewer time and overrides. Strengthen prompts, retrieval and schemas using the golden set.

17.5 Days 61 to 75: monitoring and waterfall pilot

Add one account or escrow-waterfall model and a small covenant-monitoring workflow. Keep communications as drafts. Test triggers, missing evidence, amendments, waivers and operational failure. Confirm that every calculation and proposed classification has an independent review path.

17.6 Days 76 to 90: controlled decision

Compare observed quality, time, control and incident evidence with the baseline. The authorised governance body decides whether to stop, remediate, continue shadow operation or approve a bounded production release. Any release specifies eligible transactions, users, permissions, models, tools, thresholds, sampling and rollback.




Figure 8. Ninety-day implementation roadmap

Source: Matchpoint implementation framework; AI and technology-governance context [25-29,32,35-36].

17.7 Deliverables

The minimum operating pack includes source and data inventory; entity, facility and evidence schemas; jurisdiction adapter; task and prompt register; deterministic model and test suite; golden set; access and retention design; exception and decision workflows; supplier assessment; incident and rollback plan; and an observed-benefits measurement file.

18. ROI AND OPERATING-EVIDENCE FRAMEWORK

18.1 Measurement boundary

This paper contains no approved observed Matchpoint or client productivity, cost, risk or investment-return evidence. Attributed revenue, cash cost reduction, loss reduction and alpha remain USD 0. The following measurement design and values are [Unverified illustrative scenarios] and cannot be published as achieved results.

18.2 Baseline measures

Before a pilot, measure comparable cases: elapsed time from complete intake to reviewable committee pack; analyst and specialist hours; number of source-to-model transcription steps; material exceptions found before committee; post-committee correction; covenant-calculation time; late or missed monitoring items; reviewer override; incident; and complete evidence coverage. Define case complexity and completeness so that comparisons remain meaningful.

18.3 Illustrative scenario

Assume, solely for illustration, a team processes twelve comparable underwriting cases in a quarter. The baseline requires 160 staff hours per case from intake through first committee draft. A controlled pilot records 132 hours per case, with equal or better golden-set quality and no high-severity control failure. The observed-hours difference would be 28 hours per case. It would not automatically equal cash saving because staffing, redeployment, licence, implementation, review and control costs still apply.

Assume an implementation cost of USD 180,000 and annual operating cost of USD 120,000. Assume approved loaded labour cost and observed eligible hours produce USD 210,000 of annual capacity value. The illustrative net operating value would be USD 90,000 before implementation recovery. These values are invented. They demonstrate the formula only.

Measure	Formula	Required observed evidence
hours released	comparable baseline hours minus pilot hours	time records, case scope and approved exclusions
capacity value	released hours multiplied by approved loaded rate	finance-approved rate and redeployment evidence
cash saving	actual cash cost avoided	payroll, contractor or vendor evidence
quality effect	error and exception rate change	blinded reviewed golden-set results
risk outcome	observed avoided or reduced loss	approved counterfactual and realised-loss evidence
net value	approved benefit minus implementation and operating cost	finance-approved complete cost ledger
payback	implementation cost divided by approved periodic net benefit	stable observed benefit and cost evidence

Illustrative benefits-measurement bridge




Figure 9. Illustrative benefits-measurement bridge

Source: Matchpoint measurement framework. All values are unverified illustrative assumptions.

18.4 Release gates for claims

A public claim requires approved source data, a defined baseline, comparable population, calculation, reviewer and period. Capacity, cash saving, loss avoidance and investment alpha are separate measures. A reduction in drafting time should be described as an observed drafting-time result for the measured population. It should not be converted into cash saving or loss prevention without corresponding evidence.

19. RELEASE CHECKLIST AND CLAIMS REGISTER

19.1 Underwriting release checklist

- borrower, sponsor, obligor and asset entities are resolved or explicitly conditioned;
- source population, versions, permissions and retrieval dates are recorded;
- financial periods, currencies, units and consolidation perimeter are reconciled;
- reported-to-underwriting and covenant-definition bridges are approved;
- base, downside and liquidity cases show formulas, assumptions and owners;
- facility, tranche, use-of-proceeds and repayment terms are normalised;
- security, perfection, priority, valuation and insurance evidence is recorded;
- jurisdiction lanes and bounded counsel questions are complete;
- account and escrow rules are linked to executed or approved sources;
- covenant objects, amendments, waivers and monitoring evidence are versioned;
- material conflicts, missing evidence and stale items have owners and due dates;
- conditions precedent and subsequent name satisfaction and waiver authorities;
- model, retrieval, tool and calculation versions are recorded;
- high-severity golden-set and access tests pass under the approved gate;
- committee draft distinguishes supported, represented, derived, determined and unresolved content;
- authorised reviewers approve the decision and any release or communication.

19.2 Public claims register

Claim	Evidence status in this paper	Permitted wording
Claude can support cited extraction, comparison and drafting	product capabilities documented; deployment performance requires evaluation [26-29]	Claude-assisted work within a controlled architecture
a copilot determines enforceability	unsupported	do not claim
a high evidence score predicts court outcome	unsupported	do not claim
deterministic tools can reproduce approved calculations	supported as an architecture property subject to correct implementation	versioned calculation and lineage
agents can monitor covenants without human authority	unsupported	do not claim
the workflow guarantees correct underwriting	unsupported	do not claim
the workflow prevents loss or fraud	unsupported	do not claim
T33 proves revenue, cash saving, loss reduction or alpha	no approved observed evidence supplied	attributed amounts remain USD 0
the ninety-day plan fits every organisation	unverified implementation hypothesis	adapt, test and approve locally

19.3 Model-output language

Permitted labels include extracted candidate, represented fact, deterministic calculation, supported evidence object, conditioned item, unresolved exception and specialist determination. Prohibited output includes “legally enforceable,” “fully compliant,” “no risk,” “guaranteed recovery,” “covenant waived” or “funds cleared” unless the statement comes from the properly authorised determination and is quoted within its scope.

19.4 Governance release

Production approval should identify eligible users, tasks, documents, jurisdictions, facilities, models, tools, permissions, materiality, thresholds, sampling, incident routes and expiry. Changes outside that envelope require review. The system should fail closed when required evidence, permissions or services are unavailable.

20. LIMITATIONS AND CONCLUSION

20.1 Limitations

This is a control and operating-model paper. It is not legal, regulatory, compliance, accounting, audit, tax, valuation, investment, cybersecurity, privacy or technology advice. It does not provide a full statement of UAE, DIFC, ADGM or other GCC law. Official sources can change and may have transitional or fact-specific application. Qualified specialists must determine current applicability and effect.

No live fund, borrower, facility, document population, model configuration, user interface or production result was supplied for independent testing. The worked case, weights, thresholds, times and economics are unverified illustrative scenarios. No observed evidence supports an attributed Matchpoint or client revenue increase, cash cost reduction, loss reduction or alpha result; each remains USD 0.

Language-model quality depends on source quality, retrieval, prompt, schema, model version, tools and task. Citations can be present and still fail to support a claim. Deterministic code can also contain errors. Human review can fail through time pressure, poor authority or automation bias. The architecture therefore requires layered controls, testing, sampling, incident response and an operational fallback.

20.2 Conclusion

An underwriting copilot for GCC private credit should be built as an evidence and decision-control system. The useful unit is a sourced, permissioned and reviewable credit object rather than a fluent summary. Claude can support document understanding, comparison, cited explanation and drafting. Deterministic services should control calculations, covenants, waterfalls, dates and release conditions. Jurisdiction adapters should expose evidence and legal questions while qualified counsel retains legal determinations.

The same evidence lineage can serve A3 underwriting and monitoring and A1 allocator oversight when permissions and definitions are preserved. The operating test is clear: can an authorised reviewer reproduce the conclusion, see every unresolved condition, understand who decided what, and trace the result into monitoring? A bounded pilot with observed evidence can answer that question. Public economic or risk claims should wait for approved results.

REFERENCES

- [1] Central Bank of the UAE. Credit Risk Management Standards. In force from 30 November 2024. <https://rulebook.centralbank.ae/en/rulebook/credit-risk-management-standards>
- [2] Central Bank of the UAE. Credit Risk Management Regulation. <https://rulebook.centralbank.ae/en/rulebook/credit-risk-management-regulation>
- [3] Central Bank of the UAE. Standards for Capital Adequacy of Banks in the UAE. <https://rulebook.centralbank.ae/en/rulebook/standards-capital-adequacy-banks-uae>
- [4] United Arab Emirates. Federal Decree-Law No. 50 of 2022, Promulgating the Commercial Transactions Law. <https://uaelegislation.gov.ae/en/legislations/1610>
- [5] United Arab Emirates. Federal Decree-Law No. 46 of 2021, Electronic Transactions and Trust Services. <https://uaelegislation.gov.ae/en/legislations/1539>
- [6] United Arab Emirates. Cabinet Resolution No. 28 of 2023, executive regulation concerning electronic transactions and trust services. <https://uaelegislation.gov.ae/en/legislations/2199>
- [7] United Arab Emirates. Federal Decree-Law No. 51 of 2023, Financial Reorganisation and Bankruptcy Law. <https://uaelegislation.gov.ae/en/legislations/2190>
- [8] United Arab Emirates. Cabinet Resolution No. 94 of 2024, executive regulation of the Financial Reorganisation and Bankruptcy Law. <https://uaelegislation.gov.ae/en/legislations/2582>
- [9] United Arab Emirates. Federal Decree-Law No. 35 of 2022, Law of Evidence in Civil and Commercial Transactions. <https://uaelegislation.gov.ae/en/legislations/1612>
- [10] United Arab Emirates. Federal Decree-Law No. 32 of 2021, Commercial Companies. <https://uaelegislation.gov.ae/en/legislations/1542>
- [11] United Arab Emirates. Federal Decree-Law No. 45 of 2021, Protection of Personal Data. <https://uaelegislation.gov.ae/en/legislations/1972>
- [12] United Arab Emirates. Cabinet Resolution No. 29 of 2021, executive regulation concerning securing rights in movable assets. <https://uaelegislation.gov.ae/en/legislations/1487>
- [13] Emirates Development Bank. Emirates Integrated Registries Company; Emirates Movable Collateral Registry. <https://edb.gov.ae/en/solutions/emirates-integrated-registries-company>
- [14] Government of Dubai. Law No. 8 of 2007, Concerning Escrow Accounts for Real Estate Development in the Emirate of Dubai. <https://dlp.dubai.gov.ae/Legislation%20Reference/2007/Law%20No.%20%288%29%20of%202007%20Concerning%20Escrow%20Accounts%20for%20Real%20Estate%20Development%20in%20the%20Emirate%20of%20Dubai.html>
- [15] Dubai Land Department. Frequently Asked Questions. <https://dubailand.gov.ae/en/frequently-asked-questions>
- [16] Dubai Land Department. Request for Mortgage Payment in Escrow Account. <https://dubailand.gov.ae/en/eservices/request-for-mortgage-payment-in-escrow-account/>
- [17] Dubai Financial Services Authority. Collective Investment Funds; Credit Funds. <https://www.dfsa.ae/what-we-do/collective-investment-funds>
- [18] Dubai Financial Services Authority. Legislation. <https://www.dfsa.ae/laws-rules/legal-resources/legislation>
- [19] Abu Dhabi Global Market. Rules and Regulations. <https://www.adgm.com/legal-framework/rules-and-regulations/>
- [20] Abu Dhabi Global Market Financial Services Regulatory Authority. Financial Services Regulatory Authority. <https://www.adgm.com/financial-services-regulatory-authority>
- [21] Abu Dhabi Global Market. FSRA Commences Public Consultation on Proposals to Introduce Private Credit Funds. Consultation material. <https://www.adgm.com/media/announcements/adgm-fsra-commences-public-consultation-on-proposals-to-introduce-private-credit-funds>
- [22] Dubai Financial Services Authority. DFSA Proposes Significant Updates to Its Collective Investment Fund Framework. Consultation announcement, July 2026. <https://www.dfsa.ae/news/dfs-a-proposes-significant-updates-its-collective-investment-fund-framework>
- [23] Basel Committee on Banking Supervision. Principles for the Management of Credit Risk. 30 April 2025. <https://www.bis.org/bcbs/publ/d595.htm>
- [24] IFRS Foundation. IFRS 9 Financial Instruments. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-9-financial-instruments/>
- [25] Central Bank of the UAE. Outsourcing Regulation for Banks. <https://rulebook.centralbank.ae/en/rulebook/outsourcing-regulation-banks>
- [26] Anthropic. Citations. <https://docs.anthropic.com/en/docs/build-with-claude/citations>
- [27] Anthropic. Tool Use with Claude. <https://docs.anthropic.com/en/docs/agents-and-tools/tool-use/overview>
- [28] Anthropic. Evaluating AI Systems. <https://www.anthropic.com/research/evaluating-ai-systems>
- [29] Anthropic. Building and Evaluating Trustworthy Agents. <https://www.anthropic.com/research/trustworthy-agents>
- [30] World Wide Web Consortium. PROV-O: The PROV Ontology. W3C Recommendation. <https://www.w3.org/TR/prov-o/>
- [31] Financial Stability Board. Report on Vulnerabilities in Private Credit. May 2026. <https://www.fsb.org/2026/05/report-on-vulnerabilities-in-private-credit/>
- [32] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>
- [33] International Monetary Fund. Global Financial Stability Report, April 2024. <https://www.imf.org/en/publications/gfsr/issues/2024/04/16/global-financial-stability-report-april-2024>
- [34] International Monetary Fund. The Rise and Risks of Private Credit, Chapter 2 of the April 2024 Global Financial Stability Report. <https://www.imf.org/en/-/media/files/publications/gfsr/2024/april/english/ch2.pdf>
- [35] National Institute of Standards and Technology. AI Resource Center. <https://airc.nist.gov/>
- [36] Central Bank of the UAE. Guidelines for Financial Institutions Adopting Enabling Technologies. <https://rulebook.centralbank.ae/en/rulebook/guidelines-financial-institutions-adopting-enabling-technologies>

APPENDIX A. CREDIT EVIDENCE OBJECT

Field	Entry
evidence ID	stable transaction-specific identifier
proposition	one atomic statement or value
subject and perimeter	entity, facility, asset, account or covenant
source	document ID, version, hash and exact location
evidence class	observed, represented, derived or determined
source authority	configured hierarchy for the proposition
period	valid date or measurement period
units	currency, scale and measurement basis
transformation	formula, adjustments, tool and version
permission	transaction, recipient, purpose and restrictions
conflict	linked competing objects and resolution state
reviewer	person, role, authority and date
disposition	supported, conditioned or unresolved
refresh	expiry or triggering event

APPENDIX B. COVENANT OBJECT

Field	Entry
covenant ID	stable facility-specific identifier
type	maintenance, incurrence, information, undertaking or event
source	executed clause, definitions and amendments
perimeter	facility, obligors and consolidation rules
test	formula, direction and threshold
period	test date, look-back and frequency
inputs	evidence objects and accepted transformations
adjustments	baskets, exclusions, caps and add-backs
cure	right, amount, period and treatment
grace	notice and time provisions
calculation	engine, rule version and result
variance	borrower, agent and lender comparison
interpretation	authorised determination and scope
event route	reviewer, counsel, agent and committee authority
waiver	scope, effective period and evidence
monitoring	next due date, trigger and owner

APPENDIX C. JURISDICTION ADAPTER RECORD

Field	Entry
adapter ID	jurisdiction, forum, structure and version
status	operative, consultation, superseded or archived
effective period	start, end and transition treatment
official sources	legislation, rulebook, registry and forms

Field	Entry
scope	entities, activities, assets and exclusions
terminology	approved local concepts and translations
evidence requirements	creation, perfection, priority and procedure
rule owner	legal or compliance specialist
operational owner	registry, facility or collateral function
reviewed date	specialist confirmation date
trigger	law, court, registry, document or asset change
limitation	matters requiring transaction-specific advice

APPENDIX D. MODEL RUN RECORD

Field	Entry
run ID	immutable identifier
task	classification, extraction, comparison, retrieval, explanation or draft
model and provider	exact identifier and date
prompt and schema	versioned references
source population	permitted source IDs and versions
retrieval	query, filters and returned locations
tools	names, versions, inputs and outputs
result	structured response and citations
validation	deterministic checks and failures
reviewer	identity, role and disposition
downstream use	memo, calculation, exception, report or communication
refresh	superseding evidence or configuration change