

Agentic Debt Execution in the UAE: From Lender Universe to Controlled Term-Sheet Comparison

A Governance and Execution Framework for Lender Mapping, Evidence Control, Confidentiality, Negotiation and Approval

Authored by Chennakeshav Adya, Independent Researcher

September 2026

Abstract

Debt execution involves a sequence of judgement-intensive activities: defining the financing objective, preparing evidence, identifying suitable lenders, controlling disclosure, managing questions, comparing proposals, negotiating terms and obtaining authority to sign. Agentic artificial intelligence can support parts of this sequence by retrieving information, reconciling documents, classifying lender criteria, drafting controlled work products and monitoring deadlines. It can also propagate stale information, expose confidential material, create unsupported statements, flatten material differences between proposals or act beyond delegated authority.

This paper develops a controlled agentic-debt execution framework for UAE borrowers, sponsors, advisers and lenders. It treats the agent as a bounded workflow participant rather than a decision maker. The framework establishes a mandate charter, evidence ledger, lender-fit matrix, disclosure rooms, question-and-answer control, term-sheet normalisation, approval gates and an audit record. Five original figures and five decision tables translate the framework into an execution architecture, lender-fit model, data-room readiness process, term-sheet comparison and approval-control design.

A worked example follows a hypothetical UAE operating group seeking a senior secured facility. Every amount, score, probability, timing assumption, pricing term and scenario in the example is an analytical assumption used to demonstrate the framework. Actual lender appetite, credit decisions, legal terms, regulatory treatment, pricing and execution outcomes require transaction-specific evidence and authorised professional advice. This paper does not provide legal, regulatory, tax, accounting, investment, lending, data-protection, cyber-security or technology advice.

JEL Classification: G21, G32, G34, O33, K22

Keywords: agentic AI, debt execution, UAE lending, lender mapping, term-sheet comparison, credit process, confidentiality, human oversight, debt advisory

1. Define the execution problem

Debt execution begins with a financing problem, not a technology deployment. The borrower may need acquisition finance, refinancing, working capital, project debt, a bridge, asset-backed funding or growth capital. Each objective creates different lender sets, evidence requirements, risk allocations, timing constraints and approval paths. An agent should not start searching or drafting until the purpose, amount, currency, legal perimeter, use of proceeds, target timetable and acceptable structure are documented.

The mandate charter should identify the client, adviser, authorised decision makers, conflicts process, jurisdictions, confidentiality rules, regulated-activity perimeter and permitted communications. It should state which actions the agent may perform, which outputs require review and which actions remain prohibited. Sending materials, representing facts, changing a

lender position, accepting terms and signing documents should require explicit human authority.

Success should be defined through financing fit and execution quality. Relevant measures include credible lender coverage, evidence completeness, response time, proposal comparability, retained optionality, downside protection, all-in economics and closing certainty. A fast process that circulates inaccurate information or produces an unsuitable structure damages the mandate.

2. Map the end-to-end decision chain

The execution chain can be organised into diagnosis, readiness, lender mapping, controlled outreach, diligence, proposal analysis, negotiation, approval, documentation, closing and monitoring. Each stage consumes evidence and creates a decision record. The agent can coordinate tasks across stages, while a named human owns each decision and external statement.

Dependencies require explicit sequencing. A lender universe built before the financing requirement is stable produces noise. Outreach before data-room readiness increases inconsistent answers. Pricing comparison before covenant and condition analysis understates economic differences. Legal drafting before commercial alignment creates avoidable cost. The workflow should therefore block progression when specified evidence or approvals are missing.

Parallel work remains possible within a stage. Financial reconciliation, document indexing, lender-policy research and disclosure preparation can proceed concurrently when they use the same controlled facts. The orchestration layer should expose status, owner, evidence date and blockers without inventing completion.


3. Treat the agent as a bounded participant

An agentic system can plan subtasks, call approved tools, retrieve documents and produce work products. These capabilities create operational leverage and additional risk. The design should define a narrow role for every agent: evidence retrieval, reconciliation, lender classification, drafting, comparison, workflow coordination or monitoring. Broad general-purpose authority makes responsibility difficult to trace.

The system should operate through allow-listed data sources and tools. It should receive only the information needed for the task, record the source used, preserve the original, and route uncertain or conflicting evidence to a reviewer. A lender-mapping agent should not gain authority to email lenders. A drafting agent should not amend the approved financial model. A comparison agent should not rank proposals without showing the underlying terms and weighting.

Meaningful human oversight requires time, competence and decision rights. A nominal click does not provide effective challenge when reviewers cannot inspect sources or understand material differences. Review screens should foreground changed facts, missing evidence, exceptions, sensitive data and proposed external actions.

Figure 1. Controlled agentic debt-execution architecture



Authority narrows as the workflow moves from internal analysis to external representation and binding action.

Table 1. Decision-rights matrix for an agentic debt process

Activity	Agent contribution	Required human control	Evidence retained
financing diagnosis	retrieve and structure inputs	approve objective and constraints	mandate charter and source pack
lender mapping	classify stated criteria and prior evidence	approve inclusion and prioritisation	dated lender-fit record
disclosure	assemble approved materials	approve recipient, room and release	access and transmission log
Q&A	route, retrieve and draft	fact owner and adviser approve answer	question, sources and version
term comparison	extract and normalise proposals	validate interpretation and weighting	original terms and comparison bridge
negotiation	prepare options and issue log	authorised negotiator communicates	decision and communication record
acceptance	display conditions and deviations	delegated authority approves	signed approval and final documents

Actual authority depends on the mandate, applicable law, regulation and organisational policy.

4. Establish a single evidence ledger

The evidence ledger is the controlled record of facts, documents, calculations, assumptions and open questions used in the financing. Each entry should include an owner, source, date, version, confidentiality class, permitted use and status. The ledger should distinguish management representations, audited information, adviser analysis, third-party evidence, lender statements and analytical assumptions.

Provenance matters because apparently similar values can carry different meanings. Revenue may be audited, management-reported, forecast, annualised or adjusted. Debt may be accounting debt, contractual indebtedness or a covenant definition. An agent retrieving a number should preserve

its label and context rather than convert it silently into a common field.

Conflicts should remain visible. When a model, board paper and data-room file disagree, the workflow should flag the variance, identify the relevant owners and prevent external use until resolved or expressly disclosed. Version control should show what changed and which lenders received each version.

5. Separate fact, analysis, assumption and request

External credibility depends on clear epistemic categories. A borrower fact should have an approved source. Adviser analysis should show the method and inputs. A scenario assumption should be labelled in reader-friendly language. A lender request should be recorded as the lender's position, not converted into a borrower obligation until agreed and documented.

Agents can blur these categories when they summarise. A sentence generated from several documents may combine historical evidence, management expectation and an adviser conclusion. The workflow should require sentence-level or field-level provenance for material statements and calculations. Unsupported synthesis should enter an exception queue.

The same discipline applies to absence. A missing guarantee, consent, valuation, permit or contract should remain missing. The system should request it, assign an owner and record the effect on the process. It should not fill the gap with a probable answer.

6. Create the data-room readiness gate

Readiness should be measured against the financing type and lender diligence path. A standard checklist helps, but materiality depends on the borrower, asset, security, sector, jurisdiction and use of proceeds. Core workstreams commonly include corporate, financial, commercial, debt, tax, legal, regulatory, compliance, insurance, technology, ESG and collateral evidence.

The readiness gate should test completeness, internal consistency, dates, signatures, permissions, personal data, privileged material, third-party restrictions and version status. It should also reconcile summary presentations with underlying schedules and contracts. A document being present does not mean it supports the proposed statement.

Agents can index, classify, detect duplicate files, extract dates and compare schedules. Human owners should decide relevance, privilege, disclosure and legal sufficiency. Redaction should use a reviewed policy and preserve an unredacted controlled original.

Readiness also requires a transaction-specific reconciliation pack. The historical income statement should bridge to management reporting and the financing model; cash should bridge to bank evidence; existing facilities should bridge to the debt schedule and contractual definitions; ownership should bridge across registers, licences and organisational charts. Forecast opening balances should agree with the latest historical position. Material contracts should be connected to the revenue, margin, working-capital and capital-expenditure assumptions they support. The pack should state the cut-off date and identify events after that date.

The gate can use red, amber and green status, provided the underlying issue remains visible. Green means the evidence and ownership criteria are met for the stated purpose. Amber means a disclosed limitation has an approved treatment and date. Red means the item blocks the relevant stage. A summary colour should never replace the source, issue description, materiality and

decision. This approach allows the team to progress unaffected workstreams while protecting the integrity of external disclosure.

7. Design disclosure rooms by need to know

The process should use staged disclosure rather than release every document to every prospect. An initial room may contain a teaser or lender presentation. Qualified lenders may receive a confidential information memorandum and model. Shortlisted parties may receive detailed diligence and security information. Exclusivity or documentation can justify additional access.

Each room should have an approved purpose, recipient list, confidentiality terms, expiry, download policy and document manifest. Access changes should be logged. Sensitive personal, bank, customer, employee, security, source-code and privileged information requires specific treatment under applicable requirements.

An agent may recommend a room based on stage and policy. A person should approve the recipient and release. Bulk permissions, inherited access and copied links create leakage risk; periodic entitlement review is part of execution control.

8. Define lender eligibility before searching

Eligibility criteria should reflect the proposed transaction. These can include lender type, regulated status, geography, currency, minimum and maximum ticket, tenor, amortisation, sector policy, asset eligibility, leverage tolerance, security, recourse, return, concentration, sponsor requirements, documentation standard and timing.

Mandatory exclusions should be separated from preferences. A lender unable to provide the required currency or legal lending perimeter may be ineligible. A lender with slower approval may remain eligible but score lower for a time-sensitive bridge. This distinction prevents opaque rankings from discarding viable options.

Evidence should be dated. Fund size, dry powder, mandate, team, regulatory status and portfolio appetite can change. Website language provides a lead rather than proof of current appetite. Direct lender confirmation, recent comparable transactions and approved internal evidence should carry defined weights.

The UAE market also requires attention to the legal and operating location of both borrower and lender. An onshore company, a DIFC entity, an ADGM entity and a special-purpose vehicle can present different documentation, security, enforcement, regulatory and data-handling questions. The fit record should identify the actual lending entity and booking location rather than rely on a group brand. Cross-border structures should record withholding, currency, hedging, guarantees, upstreaming, financial-assistance and perfection questions for qualified review.

Sector policy can be granular. A lender may accept technology exposure while excluding pre-revenue models, hardware concentration or particular data uses. A real-estate lender may distinguish completed income-producing assets from construction, land or operating-company risk. The taxonomy should therefore describe the actual use of proceeds, collateral, repayment source and risk mechanism. A single broad sector tag has limited value for credit fit.

9. Construct the lender universe from multiple channels

The lender universe may include UAE banks, international banks, private credit funds, asset managers, development institutions, export-credit channels, specialist finance providers, family offices, insurance capital and structured-finance investors. The relevant set depends on the transaction rather than a generic market list.

Sources can include regulatory registers, institutional websites, public transactions, fund documents, prior approved interactions and authorised adviser records. Every record should identify the entity, relevant team, evidence, date and confidence. Personal contact details should come from authorised sources and comply with applicable privacy and communication requirements.

The universe should retain rejected candidates with the reason and evidence date. This supports review and prevents repeated research. It also allows previously unsuitable lenders to be reconsidered when the structure or mandate changes.

10. Score fit transparently

A lender-fit score can prioritise work when its logic remains visible. The score should show hard eligibility, dimension weights, evidence quality, recency and material exceptions. It should not claim to predict credit approval. The output is a research and sequencing aid.


Suggested dimensions include mandate fit, ticket capacity, sector familiarity, jurisdiction, currency, tenor, leverage, collateral, return, speed, documentation, relationship and execution evidence. Weighting should reflect the borrower's objectives. A low-cost refinancing and a rescue bridge require different priorities.

Users should be able to inspect the underlying facts and change the weights. Sensitivity analysis can reveal whether the shortlist depends on one uncertain criterion. A lender with strong fit and weak evidence may require validation before outreach.

Scoring should include an explicit evidence modifier. A criterion supported by a current lender statement or approved mandate record can receive full evidential weight. A criterion inferred from a dated transaction may receive reduced weight. A criterion supported only by general marketing language should remain provisional. The modifier prevents abundant but weak public text from outranking a smaller body of reliable evidence. It also gives the research team a clear validation agenda.

The shortlist should be tested against concentration. If all preferred lenders share the same funding source, regulatory treatment, collateral view or committee calendar, the process may appear diversified while carrying one execution risk. The team can retain credible alternatives across lender type or structure where the mandate permits. Diversification should remain purposeful; adding unsuitable names creates workload and confidentiality exposure without improving optionality.

Figure 2. Evidence-weighted lender-fit funnel



Screening narrows research effort; it does not represent a lender credit decision.

Table 2. Lender-fit dimensions and controls

Dimension	Evidence	Control question	Output treatment
legal and mandate eligibility	register, mandate, direct confirmation	can the lender participate?	hard gate or exception
capacity and ticket	fund or balance-sheet evidence	can it hold the required amount?	range with evidence date
product and tenor	published criteria and transactions	does structure match the need?	fit score and caveat
sector and geography	policy, portfolio, team evidence	is exposure permitted and understood?	fit and validation request
economics	return parameters and comparables	is the likely range acceptable?	scenario, not quoted offer
execution	recent relevant closings and process	can it meet the timetable?	confidence and dependency
relationship	authorised interaction record	is access current and appropriate?	owner and next action

Weights and thresholds should be approved for the specific financing objective.

11. Validate appetite before external outreach

Research produces a hypothesis about fit. Validation establishes whether the lender is active, has capacity, can consider the jurisdiction and product, and is willing to receive information under the proposed process. This can occur through authorised relationship channels or a controlled preliminary discussion.

The validation record should state who communicated, when, what was confirmed, what remains conditional and how long the evidence remains current. Broad statements such as “interested in the UAE” should not become a confirmed ticket or pricing range. Appetite can depend on sponsor, collateral, leverage, sector concentration, internal limits and timing.

An agent can prepare questions and record approved answers. It should not impersonate a relationship owner, fabricate a contact, or treat silence as interest. Outreach requires mandate authority and an approved communication.

12. Approve the outreach sequence

Sequencing influences competition, confidentiality and workload. A broad simultaneous process may maximise coverage but strain management and increase leakage. A targeted wave can test the proposition and refine evidence. A bilateral process can protect confidentiality but reduce competitive tension. The choice should reflect the objective and risk.

The outreach plan should show lender, rationale, relationship owner, materials, release conditions, contact method, timing and stop criteria. Conflicts and restricted parties should be checked before contact. Every lender should receive consistent core facts, with any permitted tailoring documented.

The agent can monitor readiness and propose the next wave. A human should authorise each external transmission. Failed delivery, bounced contacts and out-of-office responses are operational events, not evidence of credit appetite.

13. Control the lender presentation

The lender presentation should connect purpose, business model, historical performance, forecast, leverage, cash flow, collateral, risks, mitigants and repayment. Material measures require definitions and reconciliation. Adjusted earnings, annualised revenue, backlog, pipeline and asset values should show method and limitations.

The system can compare the presentation with the financial model and evidence ledger. It should flag numbers without a source, inconsistent dates, unsupported market claims and changed definitions. Narrative generation should use approved facts and preserve management ownership of forecasts.

Risk disclosure strengthens execution when it is specific and paired with evidence or mitigation. Generic risk language does not substitute for a clear statement of customer concentration, refinancing dependency, construction risk, regulation, litigation, covenant headroom or collateral limits.

14. Route questions to accountable owners

Lender questions should enter one controlled register with the lender, date, topic, sensitivity, owner, due date, source documents, draft, approval and response version. Similar questions can be grouped, but the original wording should remain available.

The agent can retrieve relevant evidence and draft a response. The fact owner validates content; legal, compliance, tax or other specialists review where required; the transaction lead approves release. The process should show when an answer is provisional, subject to diligence or dependent on a document.

Oral responses should be captured promptly where material. If a prior answer changes, affected lenders should receive a controlled correction. The register becomes part of the execution record and supports consistent documentation.

15. Protect confidential and personal data

Debt processes can include customer contracts, employee information, bank records, identities, beneficial ownership, source of funds, security data and commercially sensitive forecasts. Data handling should follow applicable UAE, free-zone, contractual and cross-border requirements. The workflow should identify the controller, permitted purpose, access basis, retention and deletion policy.

Least-privilege access limits exposure. Agents should receive task-specific data, and external model providers should be assessed for storage, training use, sub-processors, security, audit rights, incident response and termination. Sensitive prompts and outputs should not enter uncontrolled consumer tools.

The system should detect common sensitive fields, while a person decides whether redaction or disclosure is appropriate. Automatic redaction without review can miss context or remove information needed for credit assessment.

Data minimisation should be applied at field and document level. A lender may need aggregate employee cost without names, a customer concentration schedule without initially receiving full contracts, or beneficial-ownership evidence through a controlled compliance channel. The room design should identify the minimum evidence for each stage and the trigger for deeper disclosure. This reduces exposure and can improve review by presenting relevant material deliberately.

Cross-border processing requires a documented assessment of where the data originates, where it is stored, who can access it and which contractual safeguards apply. Free-zone and federal regimes may have distinct requirements. The transaction team should obtain appropriate legal and privacy advice rather than assume one UAE rule covers every entity and recipient. The agent's connectors, telemetry, support access and backup locations belong in the same data map.

16. Maintain privilege and legal-review boundaries

Legal privilege and confidentiality depend on facts and law. The workflow should not assume that labelling a document creates privilege. Legal counsel should define handling rules, review channels and permitted recipients for legal advice, investigations, disputes and draft transaction documents.

Agents used for summarisation or drafting can create copies and metadata. The technology design should address storage, access, retention and third-party processing before privileged content is introduced. A dedicated environment may be required for sensitive work.

Commercial teams should separate business negotiation from legal interpretation. The term comparison can identify clauses and differences; qualified counsel determines legal effect and drafting sufficiency.

17. Build an auditable question-and-answer loop

An auditable loop connects each answer to the original question, approved sources, reviewers, release and subsequent changes. The record should be readable by a person without reconstructing chat history. It should retain the final external answer and material supporting evidence.

Quality controls can test numerical agreement, dates, named entities, definitions, cross-references and unresolved placeholders. An answer that passes mechanical checks still requires substantive review. The system should display confidence as evidence status rather than a precise probability unsupported by validation.

Metrics can include time to assign, time to approve, reopen rate, correction rate and unresolved critical questions. Targets should not encourage rushed answers or conceal legitimate diligence.

18. Record indications without overstatement

An indication of interest, preliminary proposal and credit-approved commitment have different evidential weight. The tracker should record form, conditions, expiry, approval status, funding source and material caveats. Verbal or informal statements should be attributed and dated.

Agents may extract proposed terms from communications, but they should preserve the original and mark missing fields. A stated margin without base rate, fees, floor, tenor, amortisation and conditions cannot support an all-in comparison.

Pipeline reporting should distinguish contacted, engaged, diligence, indicated, term sheet, credit approved, documentation and closed. Movement between stages requires defined evidence.

19. Normalise economics before ranking

Headline margin is only one component of cost. Comparison should include benchmark, floor, margin, original issue discount, upfront fee, commitment fee, utilisation fee, agency and security costs, hedging, mandatory prepayment, amortisation, cash sweep, exit fee, make-whole, warrants and other value transfer.

Cash timing matters. Fees paid at signing, interest capitalised during a grace period and a large maturity payment create different liquidity profiles. The model should calculate cash interest, total debt service, effective proceeds and expected cost under defined base and downside paths.

Tax, accounting and legal treatment require specialist review. The comparison should preserve uncertainty and avoid presenting an estimated all-in cost as a binding result.

The economics model should run at least three paths: expected utilisation, delayed utilisation and downside or early-exit. A revolving facility with an undrawn fee may be attractive under one path and expensive under another. A term loan with original issue discount reduces usable proceeds on day one. A prepayment premium can transfer material value if a refinancing or asset sale occurs. The model should therefore show nominal rate, cash payments by period, effective proceeds, total fees and residual obligation.

Currency and benchmark mechanics require separate treatment. The proposal may reference EIBOR, SOFR or another floating benchmark, include a floor, require hedging or permit a different funding currency. The team should model benchmark sensitivity, hedge cost, basis risk, settlement timing and covenant treatment using treasury-reviewed assumptions. An agent can calculate scenarios from approved inputs; it should not supply an unsupported forecast for rates or foreign exchange.

20. Compare structure and control rights

Two proposals with similar cost can allocate risk differently. Relevant terms include borrowers, guarantors, security, ranking, permitted debt, restricted payments, acquisitions, disposals, investments, financial covenants, information undertakings, events of default, cure rights, transfers, voting, amendments and remedies.

The comparison should identify the operating effect of each control. A covenant with ample opening headroom may tighten quickly under an acquisition or forecast miss. A broad material-adverse-change condition can affect execution certainty. Transfer provisions can change the future lender group.

The agent can map clauses into a taxonomy and highlight differences. The transaction team and counsel should interpret consequences and negotiate language.

Figure 3. Term-sheet normalisation bridge



Comparable economics require cash timing, structure, control rights and closing conditions in addition to headline margin.

Table 3. Normalised term-sheet comparison fields

Category	Fields	Comparison output	Required review
proceeds	commitment, availability, deductions	usable cash and timing	finance and transaction lead
interest	benchmark, floor, margin, default rate	base and downside cash interest	finance and treasury
fees	upfront, commitment, exit, make-whole	cash timing and effective cost	finance, tax and counsel
repayment	amortisation, sweep, maturity	debt-service profile	finance and board
security	assets, guarantees, ranking, recourse	claim perimeter and flexibility	counsel and credit team
covenants	tests, baskets, permissions, cure	headroom and operating constraints	management and counsel
closing	diligence, approvals, conditions	timing and execution certainty	transaction lead and counsel

The table supports commercial analysis and does not replace legal review of the actual documents.

21. Quantify covenant headroom

Covenant comparison should use a common forecast, definitions and test dates. Headroom can be measured against leverage, interest cover, debt service, liquidity, net worth, borrowing base or project ratios as applicable. The model should reconcile each lender definition rather than insert one generic ratio.

Adjustments matter. EBITDA add-backs, pro forma acquisitions, synergies, lease treatment, permitted cash netting, exceptional items and cure rights can change reported headroom. The comparison should display contractual calculation, management view and conservative analytical view where relevant.

Stress cases should show when each covenant breaches, what information must be delivered, whether a cure is available and what operational decisions become constrained. A proposal with lower pricing can be less resilient when its covenant package gives little room through the business cycle.

22. Compare execution certainty

Execution certainty includes lender authority, credit status, conditions, diligence scope, documentation burden, funding dependencies, syndication, KYC and AML, collateral perfection, third-party consents and timetable. A signed term sheet may remain subject to extensive conditions.

The system should convert conditions into an owner-and-evidence plan. It should distinguish conditions the borrower controls from lender approvals, regulatory actions or third-party dependencies. Critical-path analysis can expose a closing date unsupported by the remaining work.

Probability language should be used carefully. A historical closing rate from dissimilar transactions does not establish the chance of closing this financing. The team can use

scenario-weighted planning while recording the assumptions and avoiding false precision.

23. Design the approval-control stack

Approval should occur at several levels: use-case approval for the agent, mandate approval for the financing, evidence approval for disclosure, communication approval for outreach, commercial approval for negotiation positions and delegated authority for acceptance. Each gate should state the accountable role and required record.


The workflow should use separation of duties. The person preparing a lender comparison should not be the only person validating extracted terms. The agent configuring a recommendation should not approve its output. Sensitive permission changes and external sends can require dual control.

Emergency procedures should specify how to pause the agent, revoke access, preserve logs, notify owners and continue manually. A kill switch without an operational fallback can stop the process at the point of greatest time pressure.

Delegated authority should be represented as structured data and a readable policy. Relevant fields include legal entity, decision body, transaction type, amount, currency, term, security, permitted deviations, quorum, signatories and expiry. The system can compare a proposed action with these fields and route it to the correct authority. It should not interpret an ambiguous delegation expansively. Where authority is unclear, the workflow should stop the action and request a formal determination.

Approval records should capture the information available at the time. A later change to amount, lender, pricing, security, covenant, condition or use of proceeds may require renewed authority. The system should compare the approved proposal with the current version and highlight material deviations. Reusing an old approval after the commercial bargain changes weakens governance and can create execution and legal risk.

Figure 4. Approval-control stack for external debt execution



Every external representation and binding step has a named human authority and retained evidence.

Table 4. Stage gates and minimum evidence

Gate	Minimum evidence	Approver	Blocker examples
use case	risk assessment, tools, data map, tests	technology and control owners	prohibited data or uncontrolled tool
mandate	objective, scope, authority, conflicts	client and adviser authorities	unclear regulated perimeter or authority
readiness	reconciled model and disclosure pack	finance and transaction lead	unresolved material variance
outreach	approved lender, recipient and materials	mandate owner	restricted party or confidentiality gap
proposal	original terms and validated comparison	finance, counsel and decision team	missing economics or condition
negotiation	approved positions and fallbacks	delegated negotiator	authority or information gap
acceptance	final economics, risks, documents, approvals	delegated body or signatory	unmet condition or unresolved deviation

Requirements should be tailored to the transaction, entity and applicable control framework.

24. Define model and prompt governance

The organisation should inventory models, agents, tools, prompts, data connectors and versions used in the process. Each use case should have an owner, purpose, risk rating, approved data classes, evaluation record, monitoring plan and retirement path. Material changes should trigger review.

Prompts can encode policy and should be controlled accordingly. A prompt directing an agent to favour speed or infer missing values creates predictable risk. Approved prompt templates should require sources, evidence dates, uncertainty handling, role boundaries and escalation. User-entered instructions should not override system restrictions.

Evaluation should use representative debt-execution tasks: number extraction, definition matching, lender classification, question routing, clause comparison and sensitive-data handling. Tests should measure accuracy, source fidelity, refusal, leakage, consistency and human-review effectiveness.

25. Defend against instruction and data attacks

Documents and messages received during diligence can contain malicious or misleading instructions. An agent that treats retrieved text as commands may disclose information, change its task or call an unsafe tool. The architecture should separate data from instructions and restrict tool use through policy.

Controls include content isolation, allow-listed actions, schema validation, origin checks, output filtering and confirmation for sensitive operations. Attachments should be scanned and opened in controlled environments. Links and macros require separate handling.

The review queue should highlight attempts to request credentials, expand access, transmit files, change recipients or bypass approvals. Security testing should include realistic lender communications and data-room content.

26. Control third-party providers

The provider assessment should cover ownership, model terms, data use, retention, location, sub-processors, security, resilience, audit rights, incident notification, intellectual property, performance, change management and termination. Contractual protections should align with materiality and applicable requirements.

Provider updates can alter behaviour. The workflow should test material model, agent or connector changes before production use. Automatic updates should not silently change extraction, ranking or disclosure behaviour during a live financing.

Exit planning requires exportable data, prompts, logs, evidence links and work products. The borrower and adviser should be able to continue the transaction if the provider fails, access is suspended or risk tolerance changes.

27. Retain a complete audit trail

The audit trail should identify who initiated each task, which model and tools ran, what data was accessed, what sources supported the output, what changed, who reviewed it and whether an external action occurred. Logs should be protected against alteration and retained according to policy and applicable requirements.

Readable decision records matter more than raw event volume. A reviewer should be able to reconstruct why a lender was selected, which materials it received, how a response was approved and why a proposal was preferred. The record should include exceptions and dissent.

Retention should be purposeful. Keeping every prompt and document indefinitely can increase exposure. Legal, compliance, privacy, regulatory and client requirements should shape the schedule and deletion process.

28. Manage exceptions visibly

Exceptions include conflicting data, low-confidence extraction, missing documents, policy overrides, unusual lender terms, access failures and time-critical decisions. They should enter one queue with severity, owner, due date, decision and evidence. The system should prevent a critical unresolved item from disappearing inside a narrative summary.

Materiality rules can prioritise review, while humans retain authority to escalate. Repeated exceptions may indicate weak source data, poor prompt design, an unsuitable model or an unrealistic timetable. Root-cause analysis should feed back into the workflow.

Overrides should record the rule, rationale, authority, duration and compensating control. An undocumented manual workaround undermines the audit record.

29. Coordinate negotiation without automating authority

Negotiation involves objectives, trade-offs, sequencing, relationships and judgement. The agent can maintain the issue list, compare drafts, retrieve precedents approved for use, calculate economics and prepare options. The authorised negotiator communicates positions and interprets counterpart behaviour.

The issue list should show lender position, borrower position, rationale, value or risk effect, owner, priority, fallback and status. Economic and legal issues should be linked where they interact. A lower margin may be offered with tighter prepayment protection or control rights.

Negotiation preparation benefits from a hierarchy of objectives. The team can classify each point as essential, important, tradable or informational, then record the reason and approval. Package trades should be modelled together because concessions interact. Additional amortisation may be acceptable in exchange for a lower margin only when liquidity remains adequate. A wider acquisition basket may justify a reporting obligation. The system can surface combinations and calculate effects from approved assumptions; the negotiator chooses the position.

Communication discipline supports credibility. The team should use one authorised channel, confirm who speaks for each party, circulate accurate action notes and avoid presenting a discussion point as agreement. When several lenders participate, the process should distinguish common terms from bilateral requests and track voting implications. The agent can maintain the record, while counsel and the transaction lead determine whether language reflects the commercial understanding.

The workflow should prevent draft language or internal fallbacks from reaching the lender accidentally. External messages require recipient verification and release approval. Meeting notes should distinguish proposals, questions and agreed points.


30. Use a common value-and-risk scorecard

A proposal scorecard can combine economics, liquidity, flexibility, control, collateral, execution certainty and relationship considerations. It should show dimensions separately before applying weights. Decision makers need to see the trade-offs rather than one opaque score.

Weights should reflect the approved objective and be tested for sensitivity. A near-term acquisition may place greater weight on closing certainty and funding availability. A long-dated refinancing may value flexibility, transfer control and prepayment terms. The scorecard should allow a proposal to remain preferred for a documented strategic reason despite a lower numerical rank.

The original term sheets and comparison bridge remain authoritative evidence. The scorecard is a decision aid and should never overwrite source terms.

Figure 5. Multi-dimensional proposal decision map



Scores are illustrative; approval should consider the underlying terms, evidence and transaction objective.

Table 5. Proposal scorecard with evidence and decision ownership

Dimension	Measure	Evidence	Decision question
economics	cash cost, effective proceeds, exit cost	original terms and model	what value transfer occurs by scenario?
liquidity	draw, amortisation, sweep, maturity	cash-flow model	can the borrower fund the operating plan?
flexibility	baskets, prepayment, acquisitions	term sheet and counsel review	which strategic actions remain available?
control	covenants, information, remedies	definitions and legal analysis	when can lenders intervene?
certainty	approval, conditions, timetable	lender evidence and closing plan	what can prevent or delay funding?
relationship	capacity, behaviour, future support	authorised interaction record	how does the lender fit future needs?

Illustrative scores should be replaced by transaction-specific evidence and approved weights.

31. Construct the hypothetical mandate

Consider a hypothetical UAE operating group seeking a three-year senior secured facility equivalent to AED 300 million for refinancing and growth capital. The group has several operating subsidiaries, cross-border customers and a mixture of contracted and project revenue. The example assumes a four-month execution window and three potential lender types: a bank, a private credit fund and a club structure.

These facts are analytical assumptions for demonstrating the framework. They do not represent a client, market quotation or current lender appetite. The model assumes that financial statements, forecasts, key contracts, debt schedules and corporate records require reconciliation before outreach.

The mandate charter gives internal agents authority to index documents, reconcile approved fields, research public lender criteria and draft internal work products. External communication, disclosure, negotiation and acceptance remain subject to named human approvals.

32. Run the readiness assessment

The initial hypothetical assessment finds that the historical financial statements agree to the general ledger, while the lender presentation uses a later management perimeter. Two customer contracts support forecast revenue, one major renewal remains unsigned, and a security schedule lacks current valuations. Beneficial-ownership and corporate records are available but require a final consistency check.

The system records each variance and assigns owners. The unsigned renewal remains a forecast dependency. The missing valuation blocks any assertion about conservative collateral coverage. The process can continue with lender-universe research while external outreach stays closed.

After management approves a reconciled model and counsel validates the corporate perimeter, the readiness gate opens for a first controlled wave. The disclosure note states the renewal status and valuation timing accurately.

33. Build the hypothetical lender-fit matrix

The initial universe contains 48 institutions identified from public and approved internal sources. Mandatory eligibility removes providers without the required jurisdiction, product or ticket. Evidence-weighted scoring produces a 14-lender research set. Direct validation confirms six for the initial wave, leaves four requiring further confirmation and removes four because of current mandate or concentration constraints.

The figures are analytical assumptions. They demonstrate the distinction between database population, eligibility, research priority and current appetite. No candidate is described as committed before a formal process and lender approval.

The team reviews the six names, conflicts, relationship ownership and disclosure sequence. It approves four for first-wave outreach and retains two as a controlled second wave.

34. Control the hypothetical Q&A process

The four lenders generate questions on revenue concentration, forecast conversion, existing security, shareholder support, permitted acquisitions and collateral value. The agent groups related questions, retrieves approved evidence and routes drafts to finance, commercial, counsel and the transaction lead.

One draft incorrectly treats an unsigned renewal as contracted revenue. The provenance control flags the mismatch because the evidence ledger classifies it as forecast. The answer is corrected before release and describes the renewal timetable. Another response concerning security is held for counsel review.

The example shows the value of controls before communication. Speed comes from retrieval and routing; accuracy comes from evidence classification, accountable ownership and release authority.

35. Normalise the hypothetical proposals

Proposal A offers a lower margin and shorter availability with quarterly amortisation. Proposal B offers a higher margin, a longer draw period and a cash sweep. Proposal C proposes a club with moderate pricing, broader conditions and more complex coordination. These terms are assumptions.

The comparison converts each proposal into effective proceeds, scheduled cash debt service and scenario cost. It maps security, covenants, baskets, transfers, conditions and approval status. The headline ranking changes when upfront fees, amortisation and make-whole provisions are included.

Finance validates the calculations, counsel reviews legal interpretation and management tests operating headroom. The scorecard displays trade-offs without converting them into an automatic recommendation.

36. Test downside and closing paths

The hypothetical downside assumes delayed project conversion, lower cash collection and a postponed collateral valuation. Proposal A breaches a liquidity threshold earlier because of amortisation. Proposal B retains cash initially but restricts acquisitions. Proposal C faces a longer documentation path and coordination risk.

The closing plan maps KYC, credit approval, security, consents, documentation, conditions and funding. The critical path shows that the target date depends on corporate approvals and collateral evidence, not only lender documentation.

Decision makers receive a base and downside comparison, unresolved items and fallbacks. They can choose a preferred proposal and retain an alternative until closing certainty improves.

37. Translate the framework into an operating model

The operating model needs a mandate owner, transaction lead, finance owner, legal counsel, compliance and privacy owners, technology owner, model-risk or equivalent challenge, data-room administrator and delegated decision body. Smaller organisations can combine roles where appropriate while preserving independent review of material outputs.

Core systems include the evidence ledger, document repository, lender register, Q&A register, term comparison, issue list, approval workflow and audit log. Integrations should use least privilege and tested schemas. Manual fallback procedures should be documented.

Management information can report gate status, critical blockers, lender coverage, Q&A ageing, proposal completeness, unresolved deviations, approval status and closing path. Metrics should support decisions rather than create performative activity.

38. Implement in ninety days

Days 1 to 30 should define the mandate process, data classifications, decision rights, approved tools, evaluation cases and a narrow internal use case. A sensible starting point is document indexing and controlled term extraction using synthetic or low-sensitivity data. Baseline manual accuracy and cycle time should be measured.

Days 31 to 60 should add the evidence ledger, lender-fit research and Q&A routing. Users should test provenance, access, exception handling and human review. Security, privacy, legal and operational-resilience findings should be resolved before expanding scope.

Days 61 to 90 should pilot one controlled transaction stage, monitor errors and user behaviour, and obtain formal approval for further use. External communication should remain human-controlled. Expansion should follow evidence that the operating control works in practice.

39. Monitor performance and risk

Performance measures can include extraction accuracy, source-link completeness, reconciliation variance, review time, correction rate, false shortlist rate, question ageing and term-comparison completeness. Risk measures can include sensitive-data events, access exceptions, unsupported statements, prompt attacks, unauthorised actions and unresolved high-severity issues.

Monitoring should examine drift in data, lender criteria, models, prompts and user behaviour. A stable benchmark set can reveal regression after updates. Live incidents and near misses should influence testing and training.

Transaction-level monitoring should also test whether the workflow improves the quality of the financing process. Useful evidence includes fewer contradictory disclosures, faster resolution of material questions, more complete proposal comparisons and earlier identification of closing blockers. Activity volume, agent messages and documents processed do not establish value by themselves. The review should compare outcomes with a defined manual baseline, consider the cost of controls and investigate any change in error severity.

Control owners should receive information suited to their responsibilities. Technology teams need model and connector performance; privacy and security teams need access and incident evidence; transaction leaders need readiness and closing risk; decision makers need proposal trade-offs and unresolved exceptions. A single aggregate dashboard can conceal the detail required for effective challenge.

The governance body should decide when to restrict, pause or retire a use case. Continued use requires current evidence that benefits, controls and accountability remain acceptable.

40. Preserve judgement, evidence and authority

Agentic systems can reduce administrative friction in debt execution by retrieving evidence, coordinating work, classifying lender criteria and comparing proposals. Their value depends on a controlled operating design that preserves confidentiality, provenance, legal boundaries and delegated authority.

The framework links every stage to an owner, evidence gate and audit record. It distinguishes research from appetite, drafting from representation, comparison from interpretation and analysis from approval. These distinctions support faster work without disguising uncertainty.

The central discipline is straightforward: the system may prepare and organise; accountable people determine suitability, disclose facts, negotiate positions and commit the borrower. A well-governed process can make the execution team more informed, consistent and responsive while maintaining the standards expected in a financing transaction.

References

- [1] Central Bank of the UAE, Guidance Note on the Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning by Licensed Financial Institutions in the U.A.E., issued 11 February 2026. <https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence>
- [2] Central Bank of the UAE, Securities and Commodities Authority, Dubai Financial Services Authority and Financial Services Regulatory Authority, Guidelines for Financial Institutions Adopting Enabling Technologies, 2021. <https://rulebook.centralbank.ae/en/rulebook/guidelines-financial-institutions-adopting-enabling-technologies>
- [3] Dubai Financial Services Authority, AI Survey 2025: Generative AI adoption has nearly tripled within the DIFC as governance continues to develop, 12 November 2025. <https://www.dfsa.ae/news/new-dfsa-ai-survey-generative-ai-adoption-has-nearly-tripled-within-difc-last-12-months-governance-continues-develop>
- [4] Financial Services Regulatory Authority, Business Plan 2025-2026. <https://assets.adgm.com/download/assets/FSRA%2BBusiness%2BPlan%2B2025%2B-%2B2026.pdf/60b2624c1e7e11f093d71a0ac1bcf369>
- [5] Basel Committee on Banking Supervision, Principles for the Management of Credit Risk, 2025. <https://www.bis.org/bcbs/publ/d595.pdf>
- [6] European Banking Authority, Guidelines on loan origination and monitoring. <https://eba.europa.eu/activities/single-rulebook/regulatory-activities/credit-risk/guidelines-loan-origination-and-monitoring>
- [7] Central Bank of the UAE, Credit Risk Management Regulation. <https://rulebook.centralbank.ae/en/rulebook/credit-risk-management-regulation>
- [8] Central Bank of the UAE, Credit Risk Management Standards. <https://rulebook.centralbank.ae/en/rulebook/credit-risk-management-standards>
- [9] Central Bank of the UAE, Model Management Standards. <https://rulebook.centralbank.ae/en/rulebook/model-management-standards>
- [10] Central Bank of the UAE, Model Management Guidance. <https://rulebook.centralbank.ae/en/rulebook/model-management-guidance>
- [11] Central Bank of the UAE, Financial Stability Report 2025. https://centralbank.ae/media/p2rpfum/cbuae-fsr-report_2025_en.pdf
- [12] Central Bank of the UAE, Annual Report 2025. <https://www.centralbank.ae/media/4qbn11cl/annual-report-2025-en.pdf>
- [13] United Arab Emirates, Federal Decree by Law No. 45 of 2021 Concerning the Protection of Personal Data. <https://www.uaelegislation.gov.ae/en/legislations/1972/download>
- [14] United Arab Emirates, Federal Decree-Law No. 20 of 2018 on Anti-Money Laundering, Combating the Financing of Terrorism and Financing of Illegal Organisations, as amended. <https://uaelegislation.gov.ae/en/legislations/1014>
- [15] Central Bank of the UAE, Risk-Based Approach Guidance. <https://rulebook.centralbank.ae/en/rulebook/21-definition-and-purpose>
- [16] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, 2023. <https://doi.org/10.6028/NIST.AI.100-1>
- [17] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, 2024. <https://doi.org/10.6028/NIST.AI.600-1>
- [18] Financial Stability Board, The Financial Stability Implications of Artificial Intelligence, 2024. <https://www.fsb.org/publications/the-financial-stability-implications-of-artificial-intelligence/>
- [19] International Organization for Standardization, ISO/IEC 42001:2023 Artificial intelligence management system. <https://www.iso.org/standard/81230.html>

- [20] International Organization for Standardization, ISO/IEC 23894:2023 Artificial intelligence risk management. <https://www.iso.org/standard/77304.html>
- [21] International Organization for Standardization, ISO/IEC 27001:2022 Information security management systems. <https://www.iso.org/standard/27001>
- [22] Committee on Payments and Market Infrastructures and International Organization of Securities Commissions, Guidance on cyber resilience for financial market infrastructures. <https://www.bis.org/cpmi/publ/d146.htm>
- [23] Office of the Comptroller of the Currency, Rating Credit Risk, Comptroller's Handbook. <https://www.occ.treas.gov/publications-and-resources/publications/comptrollers-handbook/files/rating-credit-risk/pub-ch-rating-credit-risk.pdf>
- [24] Board of Governors of the Federal Reserve System, Supervisory Guidance on Model Risk Management. <https://www.federalreserve.gov/frs/guidance/supervisory-guidance-on-model-risk-management.htm>
- [25] European Banking Authority, Analysis of RegTech in the EU financial sector, 2021. https://www.eba.europa.eu/sites/default/files/document_library/Publications/Reports/2021/1015484/EBA%20analysis%20of%20RegTech%20in%20the%20EU%20financial%20sector.pdf
- [26] International Finance Corporation, Corporate Governance Methodology. <https://www.ifc.org/en/what-we-do/sector-expertise/corporate-governance/cg-methodology>

About the Author

Authored by Chennakeshav Adya, Independent Researcher

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

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

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

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