

AI-Native Investment Banking: A Control Architecture for Origination, Valuation and Execution

A Deal Operating System for Approved Data, Traceable Evidence, Human Decisions and Measurable Value

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Investment banking converts fragmented information into decisions about clients, capital, value, risk and transaction terms. Artificial intelligence can accelerate parts of this work by retrieving documents, extracting facts, identifying relationships, drafting analyses, comparing scenarios and monitoring execution. The same capabilities can amplify errors, expose confidential information, obscure source provenance and create outputs whose fluency exceeds their evidential strength. A bank or advisory firm therefore needs an operating architecture that defines where AI may be used, which data it may process, how outputs are tested and who remains accountable for each decision.

This paper develops an AI-native control architecture for origination, valuation and transaction execution. The architecture links five operating records: a deal workflow map, a data-control architecture, an evidence ledger, human approval gates and a value scorecard. It treats AI as a governed component within the deal process rather than as an independent decision maker. Each material output is connected to approved data, identifiable sources, model and prompt versions, transformation logic, reviewer evidence and a defined downstream decision.

The framework separates assistive work from judgement and authority. AI can prepare a market map, extract comparable-company data, identify diligence questions and maintain a closing tracker. The accountable professional approves client selection, conflicts clearance, valuation conclusions, advice, disclosure, negotiation positions and transaction commitments. Controls are scaled to confidentiality, regulatory exposure, financial consequence and reversibility.

All monetary amounts, productivity measures, error rates, probabilities, thresholds and timetables in this paper are hypothetical modelling inputs. They do not describe an identified institution, client, mandate, valuation opinion, securities recommendation or legal conclusion. A live implementation requires current legal, regulatory, data-protection, cyber-security, employment, intellectual-property, accounting and professional advice in every relevant jurisdiction.

JEL Classification: G24, G32, G34, G12, C55, M15, O33

Keywords: investment banking, artificial intelligence, origination, valuation, due diligence, transaction execution, model risk, evidence traceability, data governance, human oversight

1. Define AI-native investment banking as a controlled decision system

Investment banking is an evidence and judgement business. A transaction team gathers information from clients, public filings, market data, counterparties, advisers and its own prior work. It converts that information into a view on strategy, value, financing capacity, risk, process design and negotiation. The quality of the work depends on both the underlying evidence and the professional decisions applied to it.

Artificial intelligence changes the speed and scale at which information can be processed. It can search large document sets, classify text, extract structured facts, compare clauses, generate first drafts, identify anomalies and help maintain workflow records. These capabilities can reduce mechanical effort and broaden the information considered. They can also create unsupported claims, inconsistent calculations, disclosure failures and confidentiality incidents when they operate outside a controlled process.

The design question is therefore architectural. The firm must decide which work is eligible for AI assistance, which information may enter each tool, which outputs require validation, which records must be retained and which decisions remain exclusively human. NIST's AI Risk Management Framework organises AI risk activity around govern, map, measure and manage functions.[1] Its Generative AI Profile extends those functions to risks that can be novel or amplified in generative systems.[2] These functions provide a useful control vocabulary, while the transaction process supplies the operational context.

An AI-native deal system should preserve four invariants. Evidence remains identifiable. Decision authority remains assigned. Confidential information remains inside approved boundaries. Economic conclusions remain reproducible. If any invariant fails, the speed of output has little professional value.

The operating model should distinguish assistance, recommendation and authority. Assistance includes retrieval, extraction, formatting and workflow support. Recommendation includes ranked targets, scenario interpretation and suggested diligence actions. Authority includes client advice, valuation sign-off, conflicts decisions, disclosure and binding transaction commitments. The control burden increases as work moves from assistance toward authority.

2. Map the deal workflow before selecting tools

Tool selection should follow the deal process. A firm that begins with a general-purpose model can accumulate disconnected pilots without a reliable transaction record. A workflow map identifies the work products, sources, decision owners, hand-offs and approval points that must remain coherent from origination to closing.


The process begins with market and relationship intelligence. Potential clients, assets and counterparties are identified from permitted sources. Coverage professionals qualify relevance, relationship, timing and potential need. Before a mandate proceeds, conflicts, client acceptance, financial-crime and engagement controls operate according to the firm's legal and regulatory obligations.

Once engaged, the team frames the strategic question and builds an evidence room. Financial, commercial, operational, legal, technical and market information is organised. Valuation and financing analyses convert evidence into scenarios. Diligence tests the investment thesis and

exceptions. Process management coordinates outreach, confidentiality, bids, approvals, documentation, financing and closing.

At every stage, AI may assist with a bounded task. The workflow map should record the input class, permitted model environment, expected output, accountable reviewer and record-retention rule. The same task can carry different risk in different contexts. Summarising a public annual report is materially different from summarising a draft board paper containing market-sensitive information.

Figure 1. Deal workflow map



AI assistance remains bounded by approved data, evidence and human decision rights at every stage.

3. Classify information by sensitivity and permitted use

Investment-banking data cannot be treated as one undifferentiated corpus. Public filings, licensed market data, client confidential information, personal data, material non-public information and privileged legal material carry different rights and obligations. The firm should classify information before it reaches an AI system and enforce the classification through access, routing and technical controls.

Public availability does not remove all restrictions. Market-data licences can limit storage, transformation and redistribution. Website terms can restrict automated collection. Personal data can appear in public records and still require lawful, proportionate handling. Draft client materials may combine public facts with confidential assumptions and therefore inherit the higher classification.

The UAE Personal Data Protection Law establishes a federal framework for personal-data processing, security, confidentiality and cross-border transfer.[3] The UK Information Commissioner's Office states that senior management remains accountable for the data-protection implications of AI and should be able to demonstrate how data protection by design and default has been addressed.[4] The applicable legal analysis depends on the controller, processor, purpose, data subjects and jurisdictions involved.

A practical system uses a small number of enforceable classes. Each class specifies permitted sources, storage locations, model environments, retrieval indexes, retention, logging and human approvals. Default denial is appropriate when classification is absent. Users should not decide data eligibility by intuition at the moment a prompt is sent.

Table 1. Information classification and AI-use boundary

Information class	Typical examples	Permitted AI environment	Minimum control
Public and reusable	regulator publications, public filings, issuer websites	approved enterprise or isolated research environment	source link, access terms and date
Licensed	market data, research databases, transaction databases	environment expressly permitted by licence	entitlement check, usage log and redistribution control
Client confidential	management accounts, strategy, bids, forecasts	private approved environment with contractual protection	client authority, encryption, access restriction and retention rule
Personal data	employee, customer or counterparty records	approved environment after lawful-purpose assessment	minimisation, role access, transfer control and deletion
Market-sensitive	potential transactions, price-sensitive results, financing plans	segregated need-to-know environment or excluded tool	information barrier, named access and surveillance
Privileged or restricted	legal advice, regulator correspondence, specially restricted material	excluded unless counsel and policy expressly approve	privilege decision, isolated access and preserved chain of custody

The final policy requires current legal, contractual and regulatory review.

4. Build a data-control architecture around zones and policy enforcement

The technical architecture should prevent an unauthorised data path rather than rely entirely on training and reminders. A controlled design separates source systems, ingestion, governed storage, model services, output review and the permanent deal record. Policy decisions occur at each boundary.


Source systems can include public repositories, licensed databases, client virtual data rooms, customer relationship management, finance models, email and document management. A connector should operate through an approved identity and retrieve only the records required for the task. Content is scanned, classified and tagged with source, owner, date and rights before it enters a searchable index.

The model layer should expose approved models through a managed gateway. The gateway applies access rules, blocks disallowed content, records the model and version, sets retention and training controls, and captures prompts and outputs where permitted. High-risk tasks can use a private retrieval environment or a model that runs inside the controlled boundary. Vendor terms, sub-processors and data locations should be reviewed before use.

Outputs enter a review queue rather than the final client record. A reviewer checks evidence, calculations, interpretation, tone and confidentiality. Approved material is promoted into the deal workspace with its supporting ledger. Rejected material is retained or deleted according to policy, but it should inform model and process monitoring.

The architecture needs break-glass procedures. If a vendor, model or connector becomes unavailable, the team should be able to continue critical work from controlled source records. Dependency concentration and abrupt model changes can become transaction risks when a closing timetable relies on one external service.

Figure 2. Data-control architecture



Policy enforcement separates sources, model services, review and the authoritative deal record.

5. Create an evidence ledger for every material claim

Fluent output can conceal weak evidence. The evidence ledger restores traceability by connecting every material claim, extracted fact, model input and transaction conclusion to an identifiable source and review record. It should function as a working control rather than an appendix assembled after the advice has been delivered.


Each ledger entry should record the claim, source document, exact location, source date, access date, rights, extraction method, transformation, model and version, prompt or instruction, reviewer, review date and downstream use. A financial figure should also record units, currency, period, accounting basis and reconciliation. A clause extraction should retain the executed or draft document version and the relevant section.

The ledger helps distinguish four states: sourced fact, management representation, analytical assumption and professional conclusion. These states should not be blended. A management forecast can be a valid input while remaining a forecast. A model-generated peer group can support investigation while remaining subject to human selection and comparability review.

The SEC's EDGAR APIs provide real-time access to submissions and standardised XBRL facts without authentication.[5] Machine-readable sources can improve retrieval and reconciliation. The analyst still needs to test taxonomy, period, units, filing context and company-specific extensions. Automation does not remove accounting interpretation.

Evidence quality should be scored separately from claim importance. A material claim supported only by an undated presentation requires escalation. A less material operational point may need a lighter check. The ledger allows review effort to follow decision consequence.

Figure 3. Evidence ledger



Each material output retains a chain from source to decision.

Table 2. Minimum evidence-ledger fields

Ledger field	Control question	Example evidence
Claim and status	is this a fact, representation, assumption or conclusion?	stated revenue, management forecast, valuation judgement
Source and locator	can another reviewer find the same evidence?	filing URL and section, contract version and clause
Transformation	what changed between source and output?	currency conversion, normalisation, extraction or aggregation
Model record	which model, version and instruction produced assistance?	gateway log, retrieval set and prompt identifier
Review	who tested accuracy, context and consequence?	analyst check, accounting review, legal review or senior sign-off
Downstream use	which analysis or client decision depends on the claim?	comparable table, valuation range, diligence finding or negotiation term

Ledger depth should follow materiality and regulatory context.

6. Use AI in origination without turning correlation into client need

AI can expand origination coverage by combining public filings, ownership records, financing events, executive changes, strategic announcements, sector data and relationship information. It can identify patterns such as refinancing windows, portfolio concentration, capital expenditure, acquisition history or shareholder change. These patterns can help a banker prepare a relevant hypothesis.

A signal is not evidence that a company wants advice. The coverage professional should verify the trigger, understand the decision maker, test relationship context and form a specific value hypothesis. A generic message generated from scraped information can damage trust and may create data-protection, marketing or confidentiality issues.

The origination record should identify why the potential client is in scope, which public evidence supports the hypothesis, which prior relationship may be used and which contact method is permitted. AI-generated account summaries should retain citations and dates. Sensitive CRM notes should remain within the approved environment and should not be used to train a public model.

Qualification should measure likely client value and fit rather than volume of generated outreach. Useful criteria include strategic importance, transaction plausibility, decision access, mandate fit, fee potential, conflicts, capacity and timing. A human coverage owner approves the account plan and every external communication.

The system should learn from outcomes. Meetings, qualified opportunities, signed mandates, completed transactions and collected fees provide progressively stronger evidence. Open rates and generated contact counts are weak commercial measures. The value scorecard should therefore connect origination activity to verified movement through the mandate pathway.

7. Govern opportunity screening and conflicts before analysis deepens

An AI-assisted screening model can rank opportunities by mandate fit, sector relevance, transaction scale, relationship and timing. The model should expose its factors and source data. It should not create an opaque priority score that displaces coverage judgement or conflicts controls.

Conflicts analysis requires authoritative internal records, entity resolution and legal interpretation. AI can assist by identifying names, affiliates, funds, directors and prior matters. False negatives are particularly costly because a missed relationship can expose confidential information or undermine client duties. The conflicts function retains decision authority and should document the search scope, ambiguous matches and resolution.

Client acceptance can involve identity, ownership, sanctions, financial-crime, reputation, source of funds and engagement risk. AI can organise evidence and flag inconsistencies. It should not make the final acceptance decision unless a specific legal and regulatory framework permits that decision and the firm's governance has approved it. Human investigators need access to the original evidence and the basis for every alert.

The screening architecture should also prevent premature use of confidential information. A team may receive a teaser before conflicts clearance, while another part of the firm acts for a related party. Access should follow matter-level permissions and information barriers from the first document.

An opportunity becomes active only after the required approvals and engagement authority are recorded. This gate protects both the client and the firm and gives later AI activity a defined legal and operational perimeter.

8. Make valuation models reproducible and evidence-led

AI can help collect comparable-company data, map business segments, normalise disclosures, identify precedent transactions and draft scenario explanations. Valuation remains a professional conclusion. The model needs an explicit basis of value, valuation date, currency, capital structure, accounting framework, forecast and methodology.

IFRS 13 defines fair value for financial reporting as an exit price in an orderly transaction between market participants at the measurement date and requires a framework for measurement and disclosure.[6] Transaction advice can use different bases of value and negotiation contexts. The analyst should state the basis and avoid presenting a generated number as an accounting or fairness conclusion.

Comparable selection should retain inclusion and exclusion reasons. An AI system may find firms with similar language but different revenue models, geography, scale, growth, margins, capital intensity or risk. The banker should approve the peer set and record adjustments. Precedent transactions require similar care because announced headline values can omit earn-outs, assumed debt, contingent consideration and non-comparable market conditions.

Every extracted financial fact should reconcile to the source filing. Periods, units, currencies, continuing operations and accounting policy can change comparability. Formula cells should remain visible and auditable. A generated model that cannot be independently rebuilt from the ledger should not support a client conclusion.

Valuation ranges should connect to scenarios rather than imply false precision. Revenue, margin, reinvestment, discount rate, terminal assumptions and capital structure should be tested together. The client should see which variables create the range and which evidence could narrow it.

9. Control model risk, scenarios and analytical uncertainty

Investment-banking analysis combines conventional financial models with AI-assisted extraction, classification and narrative. Each component has different failure modes. A spreadsheet can contain a formula error. An extraction model can misread a table. A language model can state a plausible but unsupported conclusion. A control system should test each risk at the appropriate layer.

Validation begins with intended use. The team should define what the model is expected to do, what it must not do, the input population, performance measure, materiality threshold and escalation. Testing should include representative documents, difficult layouts, missing fields, conflicting sources, prompt variation and model updates.

FINRA's 2026 regulatory oversight report describes governance, testing, monitoring, documentation, prompt and output logs, model-version tracking and human review as relevant considerations for generative AI use.[7] FINRA's rules remain technologically neutral; supervision, communications, recordkeeping and fair-dealing obligations continue to apply when a firm uses AI.[7]

Scenario analysis should retain uncertainty rather than convert it into a single generated forecast. The model should distinguish management case, independent base case, downside and upside. Correlations and common-cause shocks should be explicit. Assumptions should be dated, owned and compared with actual outcomes.

Changes to a model, retrieval corpus or prompt can alter outputs without a visible software release. The firm should define what constitutes a material change and when revalidation is required. A closing-critical process should have a frozen configuration or documented version so that an analysis can be reproduced after the event.

10. Use AI to deepen diligence and preserve exception ownership

Diligence is an exception-finding process. AI can compare contracts, extract obligations, analyse customer cohorts, identify financial anomalies, classify technical documents and reconcile data-room indexes. It can broaden coverage and help teams focus on material inconsistencies.

The diligence plan should begin with the transaction thesis and value drivers. Each workstream defines the claim being tested, required evidence, materiality, responsible reviewer and possible deal response. AI-generated questions should be reviewed before they reach management. A large list of generic questions can consume time without improving the decision.

Document extraction should preserve version and locator. Contract terms can differ across amendments, side letters and order forms. A model that reads the base agreement without the latest amendment can produce a wrong conclusion. Financial diligence requires reconciliation among management reports, general ledger, invoices, cash, tax records and contracts.

Findings should enter an exception register with owner, evidence, economic effect, proposed response and closure standard. Possible responses include valuation adjustment, working-capital mechanism, condition precedent, representation, indemnity, covenant, remediation plan or withdrawal. AI can suggest categories; advisers approve the legal and commercial response.

The final diligence record should show both coverage and residual uncertainty. Absence of an identified issue is not proof that no issue exists. Sample limits, inaccessible records, management representations and unresolved contradictions should remain visible to decision makers.

11. Protect confidentiality, market-sensitive information and privilege

Transaction work regularly contains information whose unauthorised disclosure can affect markets, negotiations, employees and counterparties. Data leakage can occur through prompts, vendor logs, training, retrieval indexes, browser extensions, copied files or misconfigured access. The control design should therefore address the complete data path.

The SEC's amendments to Regulation S-P require covered institutions to maintain incident-response policies and procedures for unauthorised access to or use of customer information and to retain records documenting compliance.[8] The precise application depends on institution and information. The principles of written safeguards, incident response, service-provider oversight and retained evidence are relevant to any controlled deal environment.

Information barriers should operate at matter level. Users, service accounts and AI connectors receive the minimum access needed. Search indexes should not combine restricted matters into a broad corpus unless policy and technology can enforce permissions at retrieval time. Access and export logs should be reviewable.

Privilege deserves separate treatment. Sending legal advice into an unauthorised service can create legal risk. Counsel should define when AI may process privileged material, under which contractual and technical conditions and with which record. Teams should not assume that an enterprise licence resolves privilege in every jurisdiction.

Incident procedures should define detection, containment, assessment, notification, evidence preservation and client communication. A transaction team needs a rapid route to security, legal, compliance and senior management. Continuing to use a compromised system during a live

process can compound harm.

12. Separate drafting assistance from client advice and disclosure

AI can produce first drafts of presentations, information memoranda, valuation commentary, diligence summaries, process letters and document trackers. Drafting assistance can improve consistency and allow professionals to spend more time on analysis. The reviewer remains responsible for every statement released to a client or market participant.

The drafting workflow should pull approved facts from the evidence ledger and approved calculations from controlled models. Generated text should not create a new source of facts. Claims about market leadership, technology, pipeline, sustainability, synergies or performance need a reasonable basis and client confirmation.

The SEC has warned that misleading statements about AI use can violate securities laws and that public claims should have a reasonable basis and disclose company-specific risks rather than generic language.[9] The same discipline applies to transaction materials. An AI-related value claim should identify what the technology does, where it is deployed, the supporting evidence, dependencies and limitations.

Document review should include factual accuracy, numerical consistency, legal and regulatory content, confidentiality, permissions, audience and version. Cross-document checks can identify differences between a model, board paper and marketing material. The final signatory approves the document after required specialist reviews.

Draft provenance should remain proportionate. The permanent record needs the material source, transformations, reviewers and approved version. It does not need to retain every transient wording experiment where policy does not require it. The retention design should be agreed with legal, compliance and records management.

13. Coordinate execution through a single controlled transaction record

Execution involves many parties, documents, deadlines, conditions and decisions. AI can maintain action lists, compare versions, extract obligations, summarise meetings and flag dependencies. The value comes from integrating these tasks into one authoritative record.

The transaction record should identify workstream, item, owner, due date, status, evidence, blocker, decision and dependency. Meeting summaries become proposed updates until participants or the responsible owner confirm them. Email and document connectors should not automatically change closing status without approval.

Version control is critical. A term may change across bid letters, commitment papers, sale agreements and disclosure schedules. The system should identify the controlling document and effective version. Clause comparison can support counsel and bankers, while legal interpretation remains with qualified advisers.

Transaction financing requires sources-and-uses, funds flow, conditions precedent, approvals and cash timing to reconcile. An AI-generated checklist can miss a bespoke condition. The closing owner should approve the final list against executed documents and adviser confirmations.

After closing, the deal record should preserve decisions, executed documents, approvals and material evidence according to policy. A clean archive supports client service, audit, disputes, regulatory review and lessons learned. It also provides a governed base for future precedent analysis where confidentiality and rights permit reuse.

14. Design human approval gates around consequence and reversibility


Human oversight becomes meaningful when a named person has information, competence, authority and time to challenge the output. A generic requirement to review everything can become ceremonial. Approval gates should follow the consequence of error and the reversibility of the action.

Low-consequence work can use sampling and standard checks. A public-source summary used for internal orientation may need source confirmation. Medium-consequence work such as a draft comparable set requires analytical review and reconciliation. High-consequence work such as valuation advice, client disclosure, conflicts clearance, negotiation authority and binding commitment requires accountable senior approval and, where relevant, specialist review.

The gate should state entry evidence, tests, approver, decision and record. A reviewer should be able to reject, revise or escalate. Time pressure does not remove the gate; it should activate a documented exception process with defined authority and later review.

ESMA states that management bodies remain responsible for decisions when investment firms use AI and emphasises client interests, data quality, transparency and governance.[10] The context is retail investment services, while the accountability principle is useful for broader professional governance. The EU AI Act also establishes a risk-based legal framework for AI systems and requirements that depend on role and use.[11]

Figure 4. Human approval gates



Approval intensity rises with economic consequence, confidentiality and irreversibility.

Table 3. Human approval gate design

Work product	Primary AI role	Required evidence	Accountable approval
Market and company summary	retrieval and synthesis	cited public or licensed sources	coverage professional
Opportunity score	classification and ranking	factor record, relationship context and conflicts status	coverage leader
Comparable-company analysis	extraction and calculation support	source reconciliation, peer reasons and formula review	valuation lead
Diligence finding	document comparison and anomaly detection	source locator, materiality and specialist analysis	workstream lead
Client presentation or disclosure	controlled drafting	approved ledger facts, calculation checks and required specialist reviews	mandate leader
Negotiation or binding commitment	scenario support and tracking	authorised position, executed authority and current documents	authorised human decision maker

Roles and thresholds require the firm's approved governance.

15. Translate regulatory obligations into operational controls

AI governance should connect to the laws and rules that already govern the activity. A separate AI policy that is not mapped to supervision, recordkeeping, conduct, data protection, market abuse, outsourcing and cyber-security can create gaps.

FINRA Rule 3110 requires member firms to maintain a supervisory system reasonably designed for compliance, including review of investment-banking and securities business.[12] FINRA Rule 4511 addresses books and records.[13] Applicable firms should determine how prompts, outputs, approvals and communications fit their specific obligations. Other jurisdictions and business models have different rules.

The FCA describes its approach as principles-based and focused on outcomes, relying on existing frameworks including consumer duty, accountability and governance.[14] IOSCO's consultation on AI in capital markets identifies use cases, risks and policy questions across market participants.[15] The BIS and Financial Stability Board have highlighted model risk, data quality, opacity, third-party concentration, cyber risk and human oversight in financial AI.[16][17]

An obligations register should map each use case to regulated activity, client type, jurisdiction, data class, record, control owner and review cycle. Legal status can change when a tool moves from internal assistance to client-facing recommendation or automated decision. The approval process should detect that change.

Regulatory mapping should remain current. The architecture can preserve stable control functions while specific rules, guidance and model capabilities evolve. Annual policy review may be insufficient for a rapidly changing service; material model, vendor, law and use-case changes should trigger interim assessment.

16. Manage vendor, model and infrastructure dependency

AI-native investment banking can depend on model providers, cloud platforms, market-data vendors, document systems and integration services. Concentration can create correlated failure across transactions. The firm should understand which critical work relies on each provider and how it can continue during interruption.

Vendor diligence should cover security, privacy, data use, retention, sub-processors, locations, access, model updates, intellectual property, service levels, incident notification, audit rights and exit. Technical testing should confirm contractual settings in practice. A promise that client data will not train a model needs a configuration and monitoring control.

Model portability matters. Prompts and workflows can be tightly coupled to one provider's behaviour. The firm should retain test cases and output criteria so that another approved model can be evaluated. Critical calculations should remain in deterministic models or independently reproducible code rather than inside opaque generation.

NIST's Generative AI Profile identifies risks including confabulation, data privacy, information security, intellectual property, harmful bias and value-chain integration.[2] ISO/IEC 42001 specifies requirements for an AI management system and emphasises risk, opportunity, traceability and continual improvement.[18] These frameworks can support vendor and enterprise controls without replacing transaction-specific judgement.

The continuity plan should define manual alternatives, data exports, frozen copies of critical evidence, alternative communication and decision authority. A live transaction cannot wait for a vendor dispute to clarify access rights.

17. Measure value after review, rework and control cost


AI productivity should be measured at the completed work-product level. Faster first drafts can transfer work into review, correction, source checking and integration. A value scorecard should therefore include cycle time, quality, rework, control effort, vendor cost, user adoption and transaction outcome.

The baseline should use comparable tasks before implementation. Time saved should be net of prompt preparation, tool operation, checking and remediation. Quality measures can include factual exceptions, numerical differences, missed obligations, reviewer changes and client corrections. High-severity errors should carry greater weight than minor formatting changes.

The scorecard should distinguish capacity creation from cost reduction. A team may use saved time to examine more sources, strengthen client dialogue or accelerate a process. These benefits should be evidenced through completed work and client outcomes. Headcount assumptions should not be based on tool demonstrations alone.

Financial benefit can be modelled as productive hours released, avoided external cost, reduced error loss, faster mandate conversion or additional fee capacity. Each claim needs an owner and measurement period. Collected fees provide stronger evidence than pipeline estimates. Client confidentiality and conduct remain constraints even when a use case appears economically attractive.

The illustrative scorecard below uses hypothetical indices. A live firm should set baselines, thresholds and weights through controlled pilots and actual transaction evidence.



Hypothetical indices show why speed should be assessed with quality, control and commercial evidence.

18. Implement through bounded use cases and a twelve-month control roadmap

Implementation should begin with work that is useful, measurable and reversible. Public-source research, internal knowledge retrieval, document classification and controlled extraction can provide learning before the firm permits confidential drafting or transaction recommendations.

The first phase establishes governance, data classes, approved vendors, access, logging, incident response and a use-case inventory. The firm selects a small number of pilots with defined baselines and owners. Each pilot has an evidence set, test cases, acceptance thresholds and human approval path.

The second phase integrates successful pilots with controlled source systems and the evidence ledger. The team tests permissions, retrieval accuracy, versioning, record retention and continuity. Users receive role-specific training based on the actual workflow rather than generic awareness.

The third phase extends into valuation, diligence and execution only after earlier controls operate in practice. Independent assurance tests the architecture, records and permissions. Senior management reviews value, incidents, residual risk and whether each use case should continue, change or stop.

Table 4. Illustrative twelve-month implementation roadmap

Period	Operating objective	Required evidence	Release decision
Months 1-2	inventory use cases, data and vendors	process map, data classes, contract review and named owners	approve restricted pilot perimeter
Months 3-4	test public and low-sensitivity workflows	benchmark set, error analysis, review time and user feedback	approve, revise or stop each pilot
Months 5-6	deploy managed gateway and evidence ledger	access tests, logs, model record, source traceability and incident drill	permit controlled internal production
Months 7-9	integrate valuation, diligence and execution support	reconciliation, exception accuracy, specialist review and continuity test	approve defined transaction use cases
Months 10-12	validate value and independent control operation	value scorecard, audit findings, remediation and senior review	scale, constrain or retire each workflow

Timing depends on firm size, regulation, architecture and risk appetite.

19. Use audit, incidents and exceptions to improve the system

Controls should generate evidence that they operate. Audit can test whether data classification is applied, permissions match matters, sources remain traceable, models are approved, reviewers are identified and final records contain the required approvals. Sampling should include normal work and high-risk exceptions.

Monitoring should detect changes in model output, retrieval quality, error severity, user behaviour, data movement and vendor configuration. A rising acceptance rate can mean improving quality or declining reviewer challenge. Monitoring should combine quantitative signals with review of completed work products.

Incidents and near misses are learning inputs. The record should capture event, affected data and matters, detection, containment, cause, consequence, notification, remediation and control change. Individuals should be able to report concerns without incentives to conceal productivity setbacks.

An exceptions committee or equivalent authority can review urgent departures from policy. It should record the reason, scope, compensating controls, expiry and accountable owner. Permanent workarounds should not accumulate through repeated temporary exceptions.

Table 5. Control test and response matrix

Control domain	Test	Failure signal	Required response
Data boundary	attempt restricted retrieval across matters and classes	unauthorised result or export	contain access, assess exposure and repair policy enforcement
Evidence traceability	rebuild sampled claims from ledger	missing source, locator, transformation or reviewer	suspend affected use and remediate records
Model performance	rerun benchmark and difficult-document set	material accuracy or consistency decline	freeze version, investigate and revalidate
Human approval	inspect consequential outputs and authority	approval absent, late or ceremonial	escalate, retrain and redesign gate
Value	compare completed task with baseline	review cost, error or vendor cost exceeds benefit	constrain, redesign or retire workflow
Continuity	execute vendor-loss and incident exercise	critical work cannot continue or records are inaccessible	implement alternative path and retest

Test frequency and escalation follow materiality and applicable obligations.

20. Preserve professional judgement while increasing deal capacity

AI-native investment banking is credible when the firm can explain how an output was produced, which evidence supports it, who reviewed it and who authorised the resulting action. The system should make professional work more reproducible without treating professional judgement as a model output.

The five records in this paper create one operating architecture. The workflow map defines the deal process and hand-offs. The data-control architecture limits information movement and model access. The evidence ledger preserves source provenance and transformations. Human approval gates retain accountability at consequential decisions. The value scorecard tests whether the workflow creates measurable benefit after review and control effort.

The architecture can support origination, valuation, diligence and execution across jurisdictions and transaction types. Its controls remain proportional to the work. Public-source retrieval can move quickly. Confidential valuation advice and transaction authority require deeper evidence, specialist review and senior approval.

The firm should start with bounded use cases, test them against real work, preserve exception evidence and scale only after controls operate. It should measure completed outputs and client outcomes rather than generated volume. Vendor capability will continue to change, while the control objectives of confidentiality, traceability, accountability, reproducibility and client interest remain durable.

A well-governed system can allow deal teams to examine more evidence, identify exceptions earlier, maintain clearer transaction records and spend more professional time on judgement, negotiation and client decisions. The economic value arises when those improvements are verified in completed mandates and collected outcomes.

References

- [1] National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” January 2023. <https://www.nist.gov/itl/ai-risk-management-framework>
- [2] National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile,” NIST AI 600-1, July 2024. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [3] United Arab Emirates Government, “Federal Decree-Law No. 45 of 2021 Concerning the Protection of Personal Data.” <https://www.uaelegislation.gov.ae/en/legislations/1972/download>
- [4] Information Commissioner's Office, “What are the accountability and governance implications of AI?” <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/what-are-the-accountability-and-governance-implications-of-ai/>
- [5] United States Securities and Exchange Commission, “EDGAR Application Programming Interfaces.” <https://www.sec.gov/search-filings/edgar-application-programming-interfaces>
- [6] IFRS Foundation, “IFRS 13 Fair Value Measurement.” <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
- [7] Financial Industry Regulatory Authority, “GenAI: Continuing and Emerging Trends,” 2026 Annual Regulatory Oversight Report. <https://www.finra.org/rules-guidance/guidance/reports/2026-finra-annual-regulatory-oversight-report/gen-ai>
- [8] United States Securities and Exchange Commission, “Regulation S-P: Privacy of Consumer Financial Information and Safeguarding Customer Information,” May 2024. <https://www.sec.gov/rules-regulations/2024/06/s7-05-23>
- [9] United States Securities and Exchange Commission, “Chair Gary Gensler on AI Washing,” 18 March 2024. <https://www.sec.gov/newsroom/speeches-statements/sec-chair-gary-gensler-ai-washing>
- [10] European Securities and Markets Authority, “Public Statement on the use of Artificial Intelligence in the provision of retail investment services,” 30 May 2024. <https://www.esma.europa.eu/document/public-statement-ai-and-investment-services>
- [11] European Union, “Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence,” 13 June 2024. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32024R1689>
- [12] Financial Industry Regulatory Authority, “Rule 3110: Supervision.” <https://www.finra.org/rules-guidance/rulebooks/finra-rules/3110>
- [13] Financial Industry Regulatory Authority, “Rule 4511: General Requirements.” <https://www.finra.org/rules-guidance/rulebooks/finra-rules/4511>
- [14] Financial Conduct Authority, “AI and the FCA: our approach,” updated 13 February 2026. <https://www.fca.org.uk/firms/innovation/ai-approach>
- [15] International Organization of Securities Commissions, “Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges,” CR/01/2025. <https://www.iosco.org/library/pubdocs/pdf/IOSCPD788.pdf>
- [16] Bank for International Settlements, “Intelligent financial system: how AI is transforming finance,” Working Paper No. 1194, 13 June 2024. <https://www.bis.org/publ/work1194.htm>
- [17] Bank for International Settlements, “Financial stability implications of artificial intelligence: Executive Summary,” 26 June 2025. <https://www.bis.org/publications/fsi-summary-financial-stability-implications-artificial-intelligence-executive-summary>
- [18] International Organization for Standardization, “ISO/IEC 42001:2023 AI management systems.” <https://www.iso.org/standard/42001>
- [19] Dubai Financial Services Authority, “UAE regulatory authorities jointly issue Guidelines for Financial Institutions Adopting Enabling Technologies,” 15 November 2021. <https://www.dfsa.ae/news/uae-regulatory>

-authorities-jointly-issue-guidelines-financial-institutions-adopting-enabling-technologies

[20] United States Department of Justice and Federal Trade Commission, “2023 Merger Guidelines,” 18 December 2023. <https://www.justice.gov/atr/2023-merger-guidelines>

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.