

AI at the Diligence Table: Governed Analysis of Infrastructure Evidence

A decision framework for using artificial intelligence across contracts, technical records and infrastructure scenarios

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

Infrastructure diligence can require thousands of contracts, drawings, inspection records, permits, models, correspondence files and operating reports. Artificial intelligence can classify documents, retrieve clauses, compare versions, extract obligations, reconcile evidence and accelerate scenario analysis. The resulting speed creates a governance problem. A fluent answer can conceal an incomplete corpus, an obsolete document, a misread table, a missing amendment or an unsupported conclusion. A transaction team therefore needs a controlled method for deciding which AI outputs can inform triage, specialist review, valuation, financing terms and investment approval. This paper develops an Evidence-to-Decision Control Stack for governed AI analysis in infrastructure transactions. It separates source authority, corpus control, retrieval, extraction, interpretation, verification and decision rights. Every material output carries a citation to the underlying evidence, a statement of scope and uncertainty, an accountable reviewer and a defined permitted use. The framework applies to concession and offtake contracts, engineering and construction records, environmental and land evidence, operating data, financial models and investment-committee papers. The approach draws on NIST's Artificial Intelligence Risk Management Framework and Generative AI Profile, ISO/IEC 42001, the European Union Artificial Intelligence Act, secure-AI guidance from CISA and the UK National Cyber Security Centre, information-management principles associated with ISO 19650, and World Bank contract-management guidance. Three figures and three tables use hypothetical author-developed assumptions to demonstrate workflow design and control economics. They are decision aids; they do not represent vendor performance, transaction outcomes, legal conclusions or investment returns.

JEL Classification: C55, G34, G38, L74, O33

Keywords: artificial intelligence, infrastructure diligence, document analysis, contract review, technical evidence, retrieval-augmented generation, model governance, investment committee, project finance, assurance

Introduction

The infrastructure data room is rarely a coherent database. It is a negotiated archive assembled from operating systems, project folders, advisers, lenders, contractors, regulators and management. A single commercial conclusion may depend on several documents that were created for different purposes. A revenue forecast can require an offtake agreement, tariff schedule, indexation formula, dispatch history, curtailment notices and amendments. A construction conclusion can require drawings, schedules, progress certificates, variation orders, test records and correspondence. A land conclusion can require title, easements, permits, surveys and local legal interpretation.

Artificial intelligence can make this archive searchable and comparable. It can identify candidate clauses, group duplicate files, extract dates, detect inconsistencies and draft issue summaries. The same system can fail quietly. Optical character recognition can misread a value. A retrieval system can omit the controlling amendment. A language model can state a plausible answer without sufficient evidence. A user can ask a question that exceeds the corpus. A reviewer can accept a polished paragraph because its uncertainty is not visible.

The central conclusion is that AI should enter diligence as a governed evidence workflow. The system should accelerate discovery and synthesis while preserving document authority, specialist judgement and approval rights. The transaction team should know which files were searched, which passages support the answer, what was excluded, what the model did, what a person verified and which decision the output may influence.

1 Start with the decision, not the model

The design begins with a decision register. Each use case should name the decision, evidence domain, consequence of error, expected reviewer, response time and permitted reliance. “Review the contracts” is too broad. “Identify termination, change-in-control and lender-consent provisions in the executed concession suite, including amendments” is testable. “Reconcile certified construction progress with drawdown evidence through the latest reporting date” is another bounded question.

The consequence of error determines the control. A missed formatting inconsistency in an internal index has limited impact. A missed termination right can change valuation, financing and transaction structure. A mistaken safety conclusion can have physical consequences. The team should therefore assign use tiers. Discovery outputs can guide where a person looks. Analytical outputs can support specialist work when independently checked. Decision-support outputs can enter a committee paper only after defined verification. AI should not receive autonomous authority to approve a bid, certify completion, release funds or make a legal conclusion.

The decision register prevents technology-led scope expansion. A successful document search does not prove the system can interpret engineering design or forecast operating performance. Each new domain requires its own evidence, evaluation and reviewer.

2 Establish source authority

The data room should distinguish executed, approved, issued, draft, superseded and reference material. File names alone are weak evidence. The authoritative record can depend on signature, effective date, amendment sequence, issue status, document-control metadata and the governing system. Duplicate documents may differ by one clause or schedule. Scanned copies may omit attachments.

An authority register should record document identifier, title, counterparty, version, date, status, source location, confidentiality, language, governing relationship and supersession chain. The system should preserve original bytes and calculate a cryptographic hash for each ingested file. Derived text, tables and images remain linked to the original. A reviewer should be able to move from an AI statement to the exact page and source file.

The corpus should have a cut-off time. Late uploads are logged and assessed for affected outputs. Deleted or replaced files remain in the audit history. Documents received outside the formal data room should be marked by provenance. These controls matter because a correct answer against an incomplete or obsolete corpus can still be wrong for the transaction.

Table 1. Evidence classes and permitted AI use

Evidence class	Typical material	AI role	Required human control	Permitted decision use
Executed legal record	Signed concession, offtake, financing and security documents	Clause retrieval, obligation extraction and cross-reference mapping	Qualified legal review against complete document suite	Legal issue spotting and diligence support after review
Approved technical record	Issued drawings, test certificates, inspection and engineer reports	Classification, comparison and exception identification	Relevant engineer confirms document status and technical meaning	Technical diligence support and targeted follow-up
Operating record	SCADA extracts, maintenance history, incident logs and performance reports	Trend preparation, anomaly triage and reconciliation	Data owner validates completeness, units and operating context	Operating analysis after validation
Financial record	Audited accounts, management accounts, model and lender reports	Tie-out, variance analysis and scenario preparation	Finance reviewer reconciles to authoritative records	Valuation and financing support after sign-off
Management representation	Q&A, forecast, presentation and commentary	Claim extraction and evidence matching	Diligence lead tests against independent evidence	Clearly attributed, limited reliance
External reference	Regulation, market data, standards and public sources	Research, comparison and change monitoring	Date, jurisdiction and source authority checked	Context only until applied by a qualified reviewer

The classification is an author-developed control aid. Transaction teams should adapt it to the mandate, jurisdiction and professional responsibilities.

3 Build a controlled corpus

Ingestion should be reproducible. The pipeline records file receipt, virus scanning, format conversion, optical character recognition, language detection, page count, metadata extraction, segmentation and indexing. Password-protected, corrupt or unsupported files enter an exception queue. Tables, drawings and handwritten annotations may need specialist treatment. The team should measure what proportion of pages produced usable text and which pages were excluded.

Segmentation affects retrieval. A clause split from its definitions, provisos or schedule can be misunderstood. Contracts may need clause-aware chunks with parent-child links. Technical reports may need section, table and figure relationships. Correspondence may need message threads and attachments. Financial models require cell-level lineage and calculation logic; flattened text is insufficient.

The corpus manifest should be versioned. Each AI response records the manifest version, retrieval configuration, model version and prompt or workflow version. This allows the team to rerun a conclusion when new evidence arrives. It also prevents a later model output from appearing to have been available at an earlier decision date.

4 Separate retrieval from reasoning

Retrieval identifies candidate evidence. Reasoning interprets it. The distinction should remain visible. A retrieval system can miss a relevant passage because of language, terminology, scanning quality, chunk design or query formulation. A language model can then reason coherently over the wrong subset.

Material questions should use multiple retrieval paths where appropriate: exact terms, semantic search, metadata filters, cross-references and a document-family search. The result set should show excluded files and unresolved references. A contract answer should retrieve definitions, operative clause, schedules and amendments. A technical answer should retrieve measurement method, units, date, assumptions and limitations.

The output should identify whether the answer is extractive, comparative or interpretive. Extraction reports what the source states. Comparison shows differences across sources. Interpretation explains implications and requires a named reviewer. This simple distinction reduces the risk that a generated explanation is mistaken for contractual text or engineering fact.

5 Require claim-level citations

A citation should support the specific claim beside it. A paragraph-level list of files is inadequate when several statements depend on different sources. The output should provide document identifier, version, page or clause, quoted or extracted passage, and a link to the original evidence. When the evidence is tabular, the citation should include table, row, column and unit where practical.

Citation presence alone does not establish support. Evaluation should test entailment: whether the cited passage actually supports the claim. It should also test completeness: whether material qualifications and amendments were retrieved. A confidence score without these tests can be misleading. The reviewer needs the evidence, retrieval path and identified gaps.


Unsupported statements should be blocked from decision-ready outputs. The system can state that the available corpus does not answer the question. That response is valuable because it creates a precise request for information rather than a fabricated conclusion.

6 Design the Evidence-to-Decision Control Stack

The control stack has seven gates. Scope defines the decision and permitted use. Authority establishes the controlling sources. Corpus confirms ingestion and version. Retrieval assembles the evidence set. Analysis extracts and compares claims. Verification tests accuracy, completeness and specialist meaning. Decision records the human conclusion, unresolved matters and approval.

Each gate has a stop condition. The workflow stops when the decision is undefined, authoritative evidence is missing, extraction quality is inadequate, a citation does not support the claim, an amendment chain is incomplete, a specialist cannot reproduce the interpretation, or the residual uncertainty exceeds the approved tolerance. Stopping produces an exception, owner and next action.

Figure 1. The Evidence-to-Decision Control Stack



The framework is author-developed. Each gate must preserve traceability before an AI-assisted output receives greater decision authority.

7 Apply legal-document controls

Legal-document analysis should begin with document-family completeness. The executed agreement can be modified by amendments, waivers, side letters, notices and schedules. Defined terms can change meaning across the suite. Governing-law analysis and legal conclusions remain the responsibility of qualified counsel.

AI can produce a clause inventory covering term, termination, change of control, assignment, consent, step-in, force majeure, change in law, tariff, indexation, performance, liquidated damages, liability, insurance, dispute and confidentiality. The inventory should quote the operative text and identify cross-references. A comparison workflow can locate differences across concession, offtake, construction, operating and financing documents.

The control test asks whether the complete executed family was searched, whether the cited passage is current, whether definitions and qualifications were included, and whether counsel confirmed the implication. The output should separate contract wording from legal interpretation and commercial consequence.

8 Reconcile obligations across the contract suite

Infrastructure risk often sits between documents. A construction contract may promise a completion date that does not align with an offtake long-stop date. A performance standard may differ between an engineering specification and a concession requirement. Insurance obligations may not match lender expectations. A consent right can interact with security enforcement.

An obligation graph can map party, action, condition, date, dependency, remedy and source. AI can populate candidate links, while legal and commercial reviewers confirm them. The graph should retain conflicting obligations rather than selecting one automatically. It can then support a closing-condition list, consent plan, covenant matrix and integration roadmap.

Materiality should be separate from extraction. A model may identify hundreds of obligations. The transaction team decides which ones affect bid price, structure, financing, closing or operations.

9 Govern technical-document analysis

Technical evidence includes drawings, specifications, design reports, equipment lists, inspection findings, test records, maintenance history and independent-engineer reports. These documents use domain-specific terminology, revision conventions and units. Drawings and tables can contain decisive information that plain-text extraction misses.

The system should preserve page images, drawing numbers, revision status, legends, units and cross-references. Computer vision can identify candidate components or changes, but an engineer verifies interpretation. Technical analysis should distinguish design intent, installed condition, tested performance and operating result. An

approved drawing does not prove installation. A test certificate does not prove continuing performance.

Exception summaries should cite the exact record and state what further evidence is required. The workflow can accelerate navigation across large technical archives while keeping engineering authority with qualified professionals.

10 Test construction and completion evidence

Construction diligence should reconcile schedule, certified progress, variation orders, invoices, test records, permits and physical evidence. AI can compare activity descriptions, dates and amounts; it can flag progress claims without matching certification or identify recurring delay reasons. The analysis should retain the reporting cut-off and distinguish planned, reported, certified and paid progress.

Completion is a contractual and technical state. Visible work, expenditure or percentage progress may not establish it. Critical systems can remain uncommissioned. Punch-list items can restrict availability. Permits and acceptance tests can be outstanding. The independent engineer and contractual certifier remain central.

For financing, the output can support a drawdown evidence matrix. Each requested amount links to approved budget, certified work, payment status, remaining cost, contingency and unresolved exception. AI assists reconciliation; authorised people approve funding.

11 Analyse operating and maintenance records

Operating evidence can include work orders, alarms, incident reports, outage logs, condition monitoring, spare-parts records and service agreements. AI can classify events, group recurring failure modes and compare maintenance action with contractual response times. It can also reconcile narrative reports with structured operational data.

Data quality should be tested before trend interpretation. Asset identifiers, timestamps, time zones, units and status codes can change. Missing records may reflect system migration rather than improved reliability. Free-text work orders can contain copied descriptions. The workflow should show coverage and known gaps by asset and period.

An engineering reviewer determines whether an observed pattern represents degradation, operating practice or reporting behaviour. Financial analysis then links verified events to availability, cost, warranty recovery and capital expenditure.

12 Control environmental, land and permitting evidence

Environmental and land files can include licences, impact assessments, monitoring reports, title, leases, easements, rights of way, maps, correspondence and compliance notices. AI can create inventories and extract dates, conditions, reporting duties and geographic references. The legal and technical meaning depends on jurisdiction and current status.

The system should distinguish a submitted application, granted permit, renewed permit and fulfilled condition. A map should be tied to the controlling legal description and coordinate system. Environmental monitoring values should retain method, unit, detection limit, sampling location and date. A generated summary should not declare compliance when evidence is incomplete.

The output can support a permit-condition matrix and land-rights exception list. Local counsel, environmental specialists and surveyors verify conclusions.

13 Protect confidential and personal information

Infrastructure data rooms can contain commercially sensitive pricing, critical-asset layouts, employee data, security information and privileged legal material. The AI architecture should enforce access rights at document and user level. Retrieval should not expose evidence to a person who could not open the source file. Logs should record prompts, retrieval and output without creating an uncontrolled copy of sensitive material.

The deployment design should address hosting location, encryption, retention, model training, subprocessors, incident response, deletion and legal hold. Privileged material may require a separated workflow. Personal information should be processed under applicable privacy law and mandate controls.

Data minimisation improves both security and relevance. A use case should ingest only the evidence required for its decision. Vendor terms and technical architecture should be reviewed before confidential transaction material is processed.

14 Defend against prompt injection and hostile content

Retrieved documents can contain instructions, hidden text, links or manipulated content intended to influence an AI system. Even ordinary documents can include language that resembles a prompt. The system should treat retrieved content as evidence, not instructions. Tool use, external access and code execution should be restricted.

Security testing should cover prompt injection, data exfiltration, privilege escalation, malicious files, poisoned reference material and unauthorised connectors. Outputs should identify external content and prevent it from changing system policy. Sensitive actions require explicit approval outside the model.

CISA and the UK National Cyber Security Centre have published secure-AI development guidance that emphasises secure design, development, deployment and operation. Transaction workflows should apply the same lifecycle discipline to acquired and configured AI components.

15 Evaluate the complete task

Model benchmarks do not establish diligence performance. Evaluation should use representative transaction tasks and owner-controlled answer keys. The test set should include executed and draft documents, amendments, poor scans, tables, drawings, multiple languages, conflicting evidence, missing documents and questions that cannot be answered.

Metrics should cover corpus coverage, retrieval recall, citation precision, entailment, extraction accuracy, numerical accuracy, amendment handling, abstention, reviewer correction, latency and cost. Results should be stratified by document type and risk. An average score can conceal failure on the material clauses.

The test should include prospective use. Historical examples can leak into model development or prompt tuning. A protected set and independent review reduce this risk. A material model, prompt, retrieval or corpus change triggers re-evaluation.

Table 2. Minimum evaluation scorecard for AI-assisted infrastructure diligence

Dimension	Test	Evidence	Failure signal	Owner
Corpus coverage	Authoritative pages successfully ingested and searchable	Manifest, exceptions and page audit	Material file or page silently excluded	Data-room manager
Retrieval recall	Required passages returned for known questions	Protected answer set	Controlling clause or amendment absent	AI product owner
Citation support	Claim is entailed by cited passage	Independent claim review	Citation is adjacent but does not support statement	Diligence workstream lead
Numerical accuracy	Values, units, dates and formulas reproduced correctly	Source comparison and recalculation	Digit, sign, unit or period error	Finance or technical reviewer
Abstention	System declines when evidence is insufficient	Unanswerable-question set	Plausible unsupported conclusion	Model-risk owner
Specialist meaning	Qualified reviewer agrees with scoped interpretation	Legal, engineering or environmental sign-off	Correct extraction produces wrong implication	Relevant specialist
Reproducibility	Output can be regenerated from recorded versions	Run record and hashes	Conclusion changes without explained input change	Assurance lead
Decision value	Verified time or quality improvement after control cost	Time log and correction record	Review burden exceeds benefit	Mandate sponsor

Thresholds should be approved before testing and calibrated to the decision consequence. The measures below are categories, not universal performance targets.

16 Make uncertainty visible

The output should identify evidence status, not display a single opaque confidence score. Useful categories include directly extracted, calculated from cited inputs, compared across sources, specialist interpretation, management representation and unresolved. Each claim can show evidence completeness, reviewer status and date.

Uncertainty can arise from missing documents, ambiguous language, conflicting sources, poor scan quality, model limitations and future assumptions. These are different and need different responses. A missing amendment requires an information request. Ambiguous language requires counsel. Conflicting measurements require technical investigation. A forecast assumption requires scenario analysis.

The system should preserve disagreement. A decision paper can present the base conclusion, alternative interpretation, financial sensitivity and required mitigation. Compressing disagreement into a confident summary destroys information that the committee needs.

17 Preserve the human review boundary

Human review should be designed, recorded and tested. A generic instruction to “check the output” is weak. The reviewer needs a defined task: confirm document authority, compare cited passage, test amendment chain, recalculate value, assess professional meaning and record disposition.

The interface should make verification easy. It should open the exact source page beside the generated claim, show related documents and permit correction. Reviewers should see unresolved exceptions before approving an output. Material corrections feed the evaluation set and workflow improvement process.

Independence matters. The person who configured a workflow should not be the sole validator for a high-consequence use. Legal, engineering, environmental and financial conclusions remain assigned to appropriately qualified people.

18 Connect AI analysis to the financial model

Verified diligence findings should enter the model through controlled assumptions. A tariff clause can affect revenue. A completion delay can affect interest during construction. An outage pattern can affect availability and maintenance cost. A permit condition can affect capital expenditure or schedule. The evidence-to-model bridge records source, interpretation, model cell, scenario and approver.

AI can propose candidate mappings and scenario descriptions. Finance reviews formulas, units, timing and interactions. The model should distinguish observed historical data, contractual values, management forecast and author-developed scenario assumptions. Automated updates should not overwrite approved inputs without change control.


This bridge allows a committee to trace a valuation movement back to evidence. It also exposes when a narrative risk has no quantified scenario or when a model assumption lacks diligence support.

19 Use scenario analysis without fabricated precision

Scenario analysis should focus on decision variables. Examples include delay months, availability, capex overrun, tariff indexation, refinancing margin, consent timing and remediation cost. AI can help assemble evidence and calculate approved scenarios. It should not invent probability distributions or market values.

The scenario register states source, rationale, range, interaction and owner. Where evidence is limited, the model can show discrete cases and break-even points. The investment committee can then decide price, structure, condition precedent, indemnity, reserve, covenant or walk-away threshold.

Figure 2. Hypothetical authority and review curve for AI-assisted diligence



Values are hypothetical author-developed assumptions. They do not represent observed performance for any model, provider or transaction.

20 Build the investment-committee evidence pack

The committee pack should present the decision, recommendation, evidence base, unresolved matters, sensitivities, mitigations and approval requested. AI-assisted content should be traceable to verified workflow outputs. The pack should identify material reliance on management representations and documents received after the cut-off.

A claim register can link every material statement to evidence, reviewer and model impact. The committee can filter for open items, low-evidence claims and assumptions with high valuation sensitivity. This improves challenge and reduces the risk that a polished narrative hides a weak evidence chain.

The final decision record should capture conditions, dissent and residual risk. It should preserve the version of the evidence and model considered. Later learning should not rewrite the historical basis for approval.

21 Procure the system through a controlled test

Procurement should use representative data under secure conditions. Vendors process a common test set and disclose architecture, models, hosting, data use, retrieval, security, human services, subcontractors and limitations. The buyer measures outputs against the protected answer key and records manual effort.

Commercial terms should cover confidentiality, data rights, model training, service levels, audit, incident notification, material model change, business continuity, export, deletion and exit. The buyer should retain evidence and annotations needed to reproduce transaction work after the subscription ends.

The selection decision should reflect total workflow cost. A lower extraction price can be offset by poor citations, high correction rates or security remediation. A premium system can still fail when the corpus and review process are weak.

22 Measure lifecycle economics

Costs include licences, hosting, connectors, ingestion, data preparation, security, evaluation, specialist review, change control, training and assurance. Benefits can include faster navigation, broader issue coverage, reduced duplicate review, improved traceability and quicker information requests. The baseline should record current time and quality before deployment.

Value should be measured after correction and review. Generated pages are not a benefit when reviewers must rebuild them. A useful measure is verified evidence units per reviewer hour: claims, clauses, reconciliations or exceptions accepted after the defined control. Another is time to identify a material issue with complete supporting evidence.

The economic case should also recognise downside avoidance cautiously. A control can reduce exposure without proving that a loss would otherwise have occurred. Avoided-loss claims require a documented basis and should remain separate from observed efficiency.

23 Implement a stage-gated operating model

The define stage selects a bounded question and reviewer. The prepare stage creates the authority register and corpus manifest. The validate stage tests representative and adversarial cases. The pilot stage runs in advisory mode alongside the existing process. The integrate stage connects verified outputs to workstream and model controls. The scale stage expands only after observed value and stable governance.

Each gate has evidence and a stop rule. The pilot stops when material citations fail, specialists cannot reproduce conclusions, security controls are inadequate, or review cost exceeds value. A narrow use that works is preferable to a broad system that creates hidden risk.

Table 3. Stage-gated deployment roadmap

Stage	Core work	Required evidence	Decision	Stop condition
Define	Name decision, evidence, consequence and reviewer	Approved use case and authority boundary	Fund corpus preparation	No accountable decision owner
Prepare	Build source register, permissions and manifest	Traceable corpus and exception list	Fund controlled evaluation	Material authority or access gap
Validate	Test retrieval, citations, numbers, abstention and security	Protected results by risk class	Approve advisory pilot	Material task fails threshold
Pilot	Run beside existing diligence process	Reviewer corrections, latency and cost	Integrate bounded outputs	Burden or residual risk is unacceptable
Integrate	Link verified claims to workstreams and model	Sign-off, audit trail and change control	Expand to adjacent decisions	Decision lineage is incomplete
Scale	Refresh evaluation, suppliers and assurance	Stable outcomes across mandates	Renew, extend or retire	Dependency or lifecycle cost exceeds benefit

The stages are author-developed and should be adapted to the mandate, information-security requirements and professional obligations.

24 Plan the first ninety days

Days one to fifteen should select two bounded tasks, such as clause inventory and evidence-to-model reconciliation. The team defines authority, reviewers, security and success measures. Days sixteen to thirty should build the corpus manifest, protected answer set and exception process. It should include difficult scans, amendments, tables and unanswerable questions.

Days thirty-one to fifty should test candidate workflows and document failure modes. The team measures retrieval, citations, numbers, abstention, security and correction time. Days fifty-one to seventy should run an advisory pilot on a controlled workstream while the existing method remains authoritative. Reviewers record corrections and unresolved issues.

Days seventy-one to ninety should decide whether to integrate, narrow or stop. The decision pack covers observed value, residual risk, control ownership, supplier terms and the next evaluation date.

25 Recognise failure modes

Common failures include ingesting an incomplete corpus, losing revision status, treating draft language as executed, separating clauses from definitions, missing tables, accepting citations that do not support claims, and allowing a model to answer beyond the evidence. Other failures include weak permissions, confidential-data leakage, prompt injection, silent model changes and unrecorded human correction.

Governance can also become performative. A long policy does not compensate for a workflow that cannot reproduce a conclusion. A confidence score does not replace source evidence. A human in the loop does not establish control when the review task is undefined or under-resourced.

The strongest warning signal is an answer that reaches a decision faster than the evidence can be inspected. Speed is useful only when traceability and accountability move with it.


26 Implications for owners, investors and advisers

Asset owners can use the framework to organise contract and technical evidence before a financing, sale or performance review. Investors and lenders can use it to widen search while keeping material conclusions tied to verified sources. Public authorities can use it to improve consistency and auditability in complex infrastructure records, subject to procurement and legal requirements.

Advisers can connect AI-enabled evidence with transaction structure. A mandate can build the authority register, design evaluation, run the controlled workflow, reconcile findings to valuation, and prepare committee-ready evidence. Specialist counsel, engineers, environmental advisers, cybersecurity professionals and accountants retain their defined roles.

The commercial advantage lies in faster access to defensible evidence, fewer undocumented assumptions and a clearer link between diligence and decision terms.

Figure 3. Hypothetical value bridge for a governed diligence workflow



Amounts are hypothetical author-developed index points. They are not forecasts, vendor quotations or transaction-performance claims.

27 Manage multilingual and cross-border evidence

Cross-border infrastructure transactions can include contracts, permits, technical reports and correspondence in several languages. Translation quality becomes part of the evidence chain. A model can retrieve a concept expressed with different terminology, but it can also lose a qualification, defined term or jurisdiction-specific meaning. The workflow should retain the original text beside any translation and record the translation method, version and reviewer.

Material legal language should be checked by a qualified bilingual reviewer or approved legal translator. Technical terminology should be validated by a domain specialist familiar with the local system. Dates, decimal separators, currencies, units and names require explicit normalisation. Machine translation should not replace the governing-language clause or the legal rules that determine which version controls.

Retrieval evaluation should include the actual transaction languages and transliterated names. A question asked in English may need to retrieve Arabic, French, Spanish or other source material. The answer should cite the original page and identify that a translation was used. Cross-border teams should also record where data is processed, which transfer rules apply and whether export-control or sovereignty requirements restrict the workflow.

28 Create continuing assurance after closing

Diligence findings can become operating controls after acquisition or financial close. The obligation graph can feed a covenant calendar. The permit matrix can become a compliance register. The maintenance exception list can inform the first hundred days. The evidence-to-model bridge can support forecast refresh and impairment review. This continuity reduces the loss of transaction knowledge during handover.

Continuing use requires fresh governance. The operating corpus changes, users change and models change. The owner should define refresh frequency, material-change triggers, access review, incident response, evaluation cadence and retirement conditions. Outputs created for diligence should not acquire operational authority automatically. Each continuing use receives a new decision statement and validation.

Monitoring should track corpus exceptions, unsupported claims, correction rates, latency, security events and unresolved actions. Material model or supplier changes should be assessed before use. The organisation should retain the ability to export evidence, annotations and audit records. A controlled retirement process confirms that sensitive data is deleted or preserved under legal and record-management requirements.

The closing archive should identify which conclusions were verified, which assumptions remain open and which controls transfer to management. This creates a direct route from diligence to value protection without representing the transaction review as a permanent operating assurance.

29 Limitations

This paper does not establish that a specific model, vendor or workflow is fit for a transaction. Performance depends on corpus quality, document type, language, configuration, model, retrieval, security, evaluation and reviewer capacity. AI output is not legal, engineering, environmental, accounting, tax or investment advice. It does not replace the professionals responsible for those conclusions.

The regulatory position varies by jurisdiction, role and use case and can change. Standards and guidance require interpretation in context. The figures and tables use hypothetical author-developed assumptions. They do not represent measured accuracy, cost, time saving, probability, price or expected return.

An effective workflow can still miss evidence or produce an incorrect output. Material conclusions require direct source inspection, qualified review and approved decision authority.

Conclusion

Artificial intelligence can make infrastructure evidence easier to find, compare and reconcile. Its value depends on the controls that connect a generated answer to an authoritative source, a qualified reviewer and a defined decision. The Evidence-to-Decision Control Stack provides that connection.

The framework begins with the decision, establishes source authority, versions the corpus, separates retrieval from interpretation, requires claim-level citations, evaluates the complete task, preserves uncertainty and assigns decision rights. It links verified findings to models and committee papers while retaining specialist responsibility. The result is a diligence process that can gain speed and coverage without weakening the evidence chain on which infrastructure capital depends.

Frequently Asked Questions

Can AI replace legal, engineering or financial diligence reviewers?

No. AI can accelerate search, extraction, comparison and reconciliation. Qualified professionals remain responsible for legal interpretation, engineering conclusions, financial analysis and decision approval.

What is the safest first use in an infrastructure data room?

A bounded discovery task with an authoritative document list and a protected answer set is usually suitable. Clause inventory, duplicate detection and evidence requests can be tested before higher-consequence interpretation.

Does a citation make an AI answer reliable?

A citation improves traceability. A reviewer must still confirm that the cited passage supports the claim, belongs to the controlling document version and includes material definitions, qualifications and amendments.

How should the team handle a question that the data room cannot answer?

The system should state that the evidence is insufficient, identify the missing source and create a precise request for information. It should not generate a plausible conclusion from incomplete evidence.

Can AI outputs feed the valuation model automatically?

Verified findings can be mapped to controlled model assumptions. Finance should confirm the source, interpretation, formula, unit, timing and scenario before any approved input changes.

How should a vendor be evaluated?

Use representative owner-controlled evidence, known answers, amendments, difficult scans, tables, drawings, multiple languages and unanswerable questions. Measure corpus coverage, retrieval, citations, numbers, abstention, security, correction time and lifecycle cost.

What records should be retained for audit?

Retain source hashes, corpus manifest, model and workflow versions, prompts or task configuration, retrieved passages, output, reviewer corrections, approval, exceptions and the evidence considered at the decision date.

When should the workflow stop or narrow scope?

Stop or narrow it when material evidence is missing, citations fail, specialists cannot reproduce conclusions, security controls are inadequate, corrections are excessive or the complete review cost exceeds observed value.

References

1. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, 2023. <https://doi.org/10.6028/NIST.AI.100-1>
2. 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>
3. National Institute of Standards and Technology, AI Risk Management Framework Playbook. <https://airc.nist.gov/airmf-resources/playbook/>
4. National Institute of Standards and Technology, Trustworthy and Responsible AI Resource Center. <https://airc.nist.gov/>
5. National Institute of Standards and Technology, Secure Software Development Framework, SP 800-218, 2022. <https://doi.org/10.6028/NIST.SP.800-218>
6. National Institute of Standards and Technology, Cybersecurity Framework 2.0, 2024. <https://doi.org/10.6028/NIST.CSWP.29>
7. National Institute of Standards and Technology, Guide to Operational Technology Security, SP 800-82 Revision 3, 2023. <https://doi.org/10.6028/NIST.SP.800-82r3>
8. International Organization for Standardization, ISO/IEC 42001:2023 Artificial intelligence management systems. <https://www.iso.org/standard/81230.html>
9. International Organization for Standardization, ISO/IEC 23894:2023 Artificial intelligence risk management. <https://www.iso.org/standard/77304.html>
10. International Organization for Standardization, ISO/IEC 22989:2022 Artificial intelligence concepts and terminology. <https://www.iso.org/standard/74296.html>
11. International Organization for Standardization, ISO/IEC 27001:2022 Information security management systems. <https://www.iso.org/standard/27001>
12. International Organization for Standardization, ISO 31000:2018 Risk management guidelines. <https://www.iso.org/iso-31000-risk-management.html>
13. International Organization for Standardization, ISO 19650-1:2018 Information management using building information modelling. <https://www.iso.org/standard/68078.html>
14. UK BIM Framework, Information Management according to BS EN ISO 19650: Guidance Part 1, 2019. https://ukbimframework.org/wp-content/uploads/2019/10/Information-Management-according-to-BS-EN-ISO-19650_-_Guidance-Part-1_Concepts_2ndEdition.pdf
15. UK BIM Framework, Guidance Part C: Facilitating the Common Data Environment Workflow, 2020. https://ukbimframework.org/wp-content/uploads/2020/09/Guidance-Part-C_Facilitating-the-common-data-environment-workflow-and-technical-solutions_Edition-1.pdf
16. European Parliament and Council, Regulation (EU) 2024/1689, Artificial Intelligence Act. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
17. European Data Protection Board, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models. https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en
18. European Union Agency for Cybersecurity, Multilayer Framework for Good Cybersecurity Practices for AI, 2023. <https://www.enisa.europa.eu/publications/multilayer-framework-for-good-cybersecurity-practices-for-ai>
19. Cybersecurity and Infrastructure Security Agency and UK National Cyber Security Centre, Guidelines for Secure AI System Development, 2023. <https://www.cisa.gov/news-events/alerts/2023/11/26/cisa-and-uk-ncsc-unveil-joint-guidelines-secure-ai-system-development>
20. Cybersecurity and Infrastructure Security Agency, Roadmap for Artificial Intelligence 2023-2024. https://www.cisa.gov/sites/default/files/2023-11/2023-2024_CISA-Roadmap-for-AI_508c.pdf
21. UK National Cyber Security Centre, Guidelines for Secure AI System Development, 2023. <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
22. UK National Cyber Security Centre, Guidelines for Secure AI System Development: Deploying AI Systems Securely, 2024. <https://www.ncsc.gov.uk/guidance/deploying-ai-systems-securely>
23. World Bank, Contract Management Practice Procurement Guidance, 2024. <https://thedocs.worldbank.org/en/doc/a5487590cccc42b5709816f40ae8b068-0290012024/contract-management-practice-procurement-guidance-june-2024-final>
24. World Bank, Systematic Tracking of Exchanges in Procurement. <https://www.worldbank.org/en/projects-operations/products-and-services/brief/systematic-tracking-of-exchanges-in-procurement-step>
25. World Bank, STEP Contract Management Module, 2022. <https://www.worldbank.org/en/news/factsheet/2022/08/24/step-contract-management-module>
26. World Bank, Procurement Regulations for IPF Borrowers, Seventh Edition, 2025. <https://thedocs.worldbank.org/en/doc/c84273d1b230aeb2b0b8134de5dc8cd7-0290012025/original/Procurement-Regulations-7th-Edition-Sep-2025.pdf>
27. World Bank PPP Resource Center, Managing PPP Contracts after Financial Close, 2018. <https://ppp.worldbank.org/library/managing-ppp-contracts-after-financial-close>
28. Global Infrastructure Hub, PPP Contract Management Tool. <https://managingppp.github.org/>
29. OECD, Recommendation of the Council on Artificial Intelligence, amended 2024. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
30. OECD, Due Diligence Guidance for Responsible Business Conduct, 2018. <https://mneguidelines.oecd.org/OECD-Due-Diligence-Guidance-for-Responsible-Business-Conduct.pdf>
31. International Finance Corporation, Performance Standards on Environmental and Social Sustainability, 2012. <https://www.ifc.org/en/insights-reports/2012/ifc-performance-standards>
32. Equator Principles Association, The Equator Principles, EP4, 2020. https://equator-principles.com/app/uploads/The-Equator-Principles_EP4_July2020.pdf
33. Financial Action Task Force, Guidance on Digital Identity, 2020. <https://www.fatf-gafi.org/en/publications/FinancialInclusionandnpoissues/Digital-identity-guidance.html>
34. Bank for International Settlements, Principles for Operational Resilience, 2021. <https://www.bis.org/bcb/publ/d516.htm>
35. IFRS Foundation, IFRS 13 Fair Value Measurement. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
36. IFRS Foundation, IAS 36 Impairment of Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>

37. IFRS Foundation, IAS 37 Provisions, Contingent Liabilities and Contingent Assets.
<https://www.ifrs.org/issued-standards/list-of-standards/ias-37-provisions-contingent-liabilities-and-contingent-assets/>
38. Project Management Institute, A Guide to the Project Management Body of Knowledge, Seventh Edition, 2021.
<https://www.pmi.org/pmbok-guide-standards/foundational/pmbok>
39. Association for the Advancement of Cost Engineering International, Total Cost Management Framework.
<https://web.aacei.org/resources/publications/tcm>
40. International Association for Contract and Commercial Management, Contract and Commercial Management resources.
<https://www.worldcc.com/Resources>
41. Open Web Application Security Project, OWASP Top 10 for Large Language Model Applications. <https://genai.owasp.org/llm-top-10/>
42. MITRE, ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems. <https://atlas.mitre.org/>
43. Cloud Security Alliance, Security Guidance for Critical Areas of Focus in Cloud Computing.
<https://cloudsecurityalliance.org/artifacts/security-guidance-v5>
44. Information Commissioner's Office, Guidance on AI and Data Protection.
<https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/>
45. UK Government, AI Playbook for the UK Government, 2025.
<https://www.gov.uk/government/publications/ai-playbook-for-the-uk-government>
46. United States Government Accountability Office, Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities, 2021. <https://www.gao.gov/products/gao-21-519sp>
47. United States Office of Management and Budget, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence, M-24-10, 2024. <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>
48. United Nations Educational, Scientific and Cultural Organization, Recommendation on the Ethics of Artificial Intelligence, 2021.
<https://unesdoc.unesco.org/ark:/48223/pf0000381137>

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.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.