AI in Engineering Diligence: Detecting Schedule, Cost and Reliability Risk before Signing

An Evidence-Control Framework for Project Data, Anomaly Triage and Transaction Decisions

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Engineering diligence often arrives late in a transaction and operates across fragmented project controls, maintenance systems, design records, operational technology, contracts and management presentations. The resulting evidence can be too large for manual review and too inconsistent for direct comparison. Artificial intelligence can accelerate document classification, data reconciliation, schedule testing, anomaly detection and reliability analysis. Its transaction value depends on controlled inputs, engineering causality, reproducible calculations and accountable human review.

This paper develops an evidence-control framework for using AI in engineering diligence before signing. It connects five records: a project-data map, anomaly-triage queue, schedule-risk heat map, reliability curve and diligence decision tree. The framework begins with the deal thesis and engineering perimeter, then maps source systems, identifiers, units, versions and access rights. It normalises the work breakdown structure, cost accounts, schedule activities, equipment hierarchy and maintenance records before analytical models are applied.

The approach separates discovery from conclusion. Automated tests can identify broken schedule logic, improbable progress, cost-to-complete inconsistency, repeated failure modes, deferred maintenance, weak data provenance and operational-technology exposure. A qualified reviewer then tests materiality, causality, alternative explanations and transaction relevance. Confirmed findings are linked to valuation, debt capacity, purchase-price mechanics, conditions, warranties, indemnities, integration priorities and governance after completion.

All amounts, probabilities, scores, thresholds, durations, cost movements and reliability assumptions in this paper are hypothetical modelling inputs. They do not describe an identified target, project, seller, buyer, financing, valuation, investment recommendation or legal conclusion. A live transaction requires current engineering, commercial, financial, legal, tax, regulatory, environmental, insurance, cyber and accounting diligence by qualified advisers.

JEL Classification: G34, G32, L64, C53, D81, M15, O32

Keywords: engineering diligence, artificial intelligence, project schedule, cost risk, reliability, maintenance, operational technology, M&A, transaction diligence, anomaly detection

1. Use AI to widen the evidence field while preserving decision authority

Engineering diligence is an evidence-conversion exercise. The buyer receives drawings, schedules, cost reports, maintenance histories, inspection records, contracts, operating data and management explanations. Each source describes a different part of the asset or project. The transaction team must determine which records are current, whether they reconcile and how the identified conditions affect value, cash, completion and continuity.

AI can expand the amount of material reviewed within the transaction timetable. Document models can classify files, extract identifiers and connect revisions. Analytical models can scan schedule networks, cost movements and maintenance histories. Statistical methods can surface unusual failure clusters, repeated overrides or inconsistent progress. These capabilities improve coverage when the evidence is controlled. They create risk when a generated summary loses provenance or an anomaly score is treated as an engineering conclusion.

The operating principle is controlled augmentation. The system should show the source, transformation, model version, output and reviewer disposition for every material signal. The engineering reviewer retains responsibility for causality and technical judgement. The deal team retains responsibility for valuation, structure and negotiation. The board or investment committee retains its decision rights.

NIST's AI Risk Management Framework organises AI risk through govern, map, measure and manage functions.[3] Its Generative AI Profile adds actions for risks specific to generative systems.[4] ISO/IEC 42001 provides an AI management-system structure covering policies, objectives, processes and continuous improvement.[6] These frameworks support a diligence environment in which speed, traceability, security and accountability are designed together.

The first practical output should be an analysis register. It records the question, evidence population, method, limitations, reviewer, materiality threshold and permitted transaction use. A model that can cluster maintenance narratives may support discovery. It should not independently determine remaining useful life, safety compliance or purchase price. The approved use should be explicit before the model touches confidential target data.

2. Fix the transaction thesis and engineering perimeter

The diligence scope should start from the value thesis. A buyer acquiring an engineering contractor may care about backlog quality, cost-to-complete, claims, equipment capacity and key-person dependence. A buyer acquiring an industrial asset may care about integrity, reliability, maintenance capital expenditure, permits and operating continuity. A project investor may focus on completion, performance testing, interface risk and liquidated damages.

The engineering perimeter records the legal entities, sites, projects, production lines, utilities, systems, contracts and periods included in the analysis. It also identifies exclusions. A group may share maintenance resources, engineering software, laboratories, spares, control rooms and intellectual property across sites. The diligence team should distinguish assets transferred at completion from services that require a transition arrangement.

The time perimeter matters. A twelve-month maintenance extract may miss a multi-year turnaround cycle. A current schedule may conceal earlier re-baselines. A recent reliability improvement may follow an abnormal outage period. The evidence plan should cover a period long enough to observe representative operating, construction and maintenance states.

The decision perimeter translates the deal thesis into testable questions. Can the current programme reach mechanical completion within the seller's forecast? Is the remaining contingency adequate for known and latent scope? Can the installed asset achieve the output and availability assumed in the valuation? Which findings require price, escrow, retention, a closing condition or a funded ownership action?

Scope control prevents analytical drift. The team should approve additions when evidence points to a new material risk. It should also record questions deferred to confirmatory diligence, specialist advice or post-completion verification. A visible scope log preserves the distinction between incomplete evidence and an adverse conclusion.

3. Build a project-data map before running models

The project-data map is the foundation of the analysis. It inventories each source, owner, system, extraction date, period, unit, identifier, revision, access restriction and transaction purpose. It covers structured records such as schedule activities and work orders, semi-structured records such as inspection sheets, and unstructured records such as progress narratives, correspondence and photographs.

Figure 1. Project-data map from source systems to transaction decisions

EACH MATERIAL SIGNAL SHOULD RETAIN A TRACEABLE PATH TO SOURCE EVIDENCE

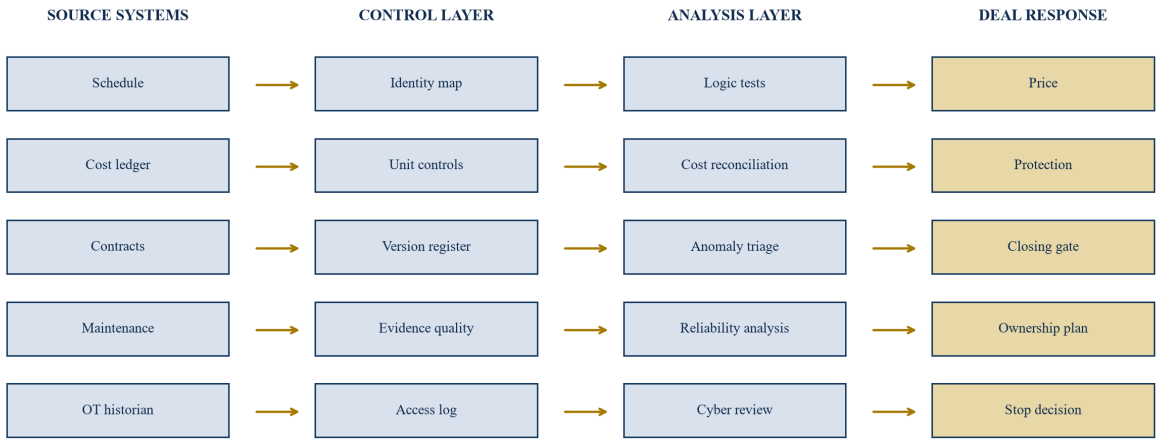

The diagram shows a controlled analytical path; actual systems and decision rights depend on the transaction.

Identifiers create the first reconciliation challenge. A pump may have a tag in the asset register, another in the maintenance system and a third in the historian. A project activity may use a package code that differs from the cost account. The identity map should preserve original identifiers and create controlled crosswalks rather than overwrite source values.

Provenance should survive every transformation. If a model converts free-text work orders into failure categories, the resulting record should link to the original text and transformation rule. If dates are standardised, the original time zone and timestamp should remain available. If duplicate records are removed, the duplicate rule and retained record should be documented.

ISO 55001:2024 strengthens the emphasis on decision-making, value, risk, data and predictive action in asset management.[7] ISO 55013 addresses the management of data assets in support of asset-management objectives.[8] In transaction diligence, these principles translate into a data room that can support both analysis before signing and controlled operation after completion.

Table 1. Engineering diligence data and control register

Data domain	Primary record	Common inconsistency	Required control
Project schedule	native schedule file, calendars and baselines	PDF differs from native logic	file hash, version and baseline register
Cost	ledger, commitments and forecast	cost account does not map to work package	controlled account crosswalk
Asset	equipment hierarchy and design duty	duplicate or missing tag	identity map and physical verification
Maintenance	work orders, inspections and backlog	free text masks repeated failure	taxonomy with source-text retention
Operations	historian, alarms and production	time, unit or sensor mismatch	unit dictionary and time alignment

Required controls depend on asset type, deal perimeter and materiality.

4. Normalise the work, cost and equipment structures

Cross-system analysis requires a common analytical grain. Schedule activities, cost accounts, contracts, equipment tags and maintenance work orders rarely share one hierarchy. The diligence team should establish a controlled map that connects each item to site, system, work package, responsible party and transaction question.

The work breakdown structure should represent deliverables and scope. The organisation breakdown structure shows responsibility. The cost breakdown structure records expenditure. The equipment hierarchy shows physical parent-child relationships. Their purposes differ, so the analytical model should connect them without collapsing them into a single artificial hierarchy.

Units and calendars need equal discipline. Cost data may mix currencies, nominal and real values, and tax treatments. Schedule data may use different working weeks, holidays and shifts. Reliability data may express runtime in hours, cycles, starts or production tonnes. The normalisation layer should record each conversion and retain source values.

Version control is essential because transaction evidence changes during diligence. A revised completion forecast may incorporate a new supplier date. A new maintenance extract may include work orders closed after the valuation date. Each analysis run should identify the source snapshot and effective date. Material revisions should be compared with prior versions to show which assumptions or records changed.

The output is a reconciliation table that shows unmatched records, one-to-many relationships and unresolved ambiguities. These are findings in their own right. A cost account with no schedule activity can represent unplanned scope. A critical schedule activity with no committed cost can indicate an incomplete forecast. An equipment tag with failures but no asset-register record can indicate weak control over the installed base.

5. Create an anomaly queue rather than a conclusion engine

An anomaly is a signal that deserves investigation. It can arise from a rule, statistical model, graph test, language model or engineering threshold. The system should create a queue with the signal, source records, method, confidence, potential materiality and reviewer disposition.

False positives are expected. A zero-duration milestone may be valid. A cost spike may reflect a planned mobilisation. Repeated equipment alarms may arise from a sensor fault rather than physical degradation. Human review should test the operating context and alternative explanations before a signal becomes a confirmed issue.

False negatives also matter. A model trained on historical maintenance categories may miss a new failure mode. Sparse failure data may make an unreliable asset look stable. Management may have changed coding practices after a reorganisation. The diligence design should combine several tests and retain manual sampling for high-consequence records.

Figure 2. Anomaly-triage funnel from machine signal to deal action

AUTOMATION EXPANDS COVERAGE; CONTROLLED REVIEW CREATES TRANSACTION EVIDENCE

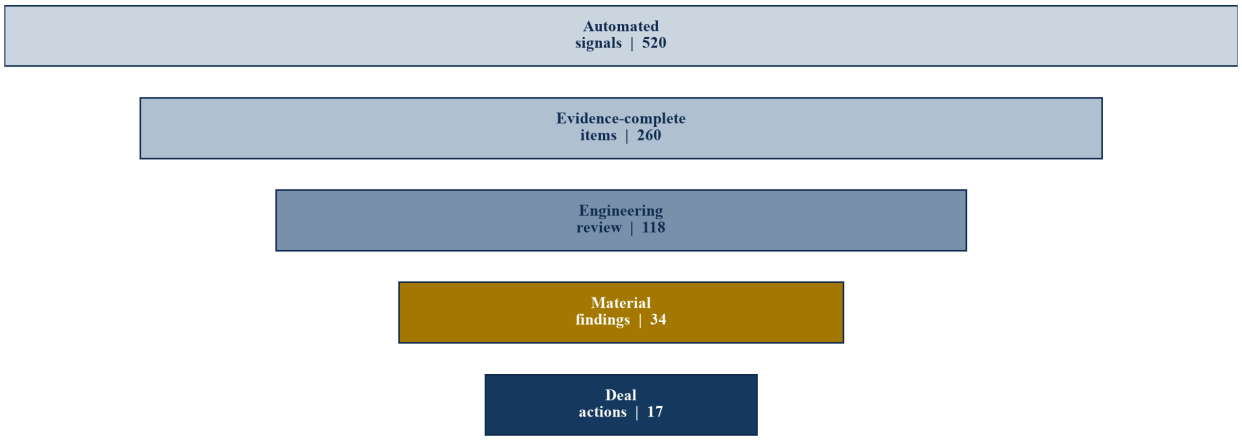

Hypothetical volumes illustrate workflow discipline; they do not describe a transaction.

Table 2. Anomaly class, validation and permitted response

Signal class	Example	Validation step	Permitted initial use
Logic	activity without predecessor	inspect network and planning basis	schedule question
Financial	forecast falls while scope grows	reconcile scope, commitments and contingency	cost challenge
Reliability	repeated corrective work on one tag	inspect failure mode, duty and maintenance	specialist review
Operational	unexplained set-point overrides	test event history and operator procedure	control-risk review
Documentary	conflicting drawing revision	confirm approved-for-use status	configuration finding

A model score alone should not determine a transaction conclusion.

The queue should record closure evidence. A signal can be confirmed, rejected, merged, deferred or left unresolved. Rejected signals should remain in the audit trail because repeated false positives can reveal a weak model or data problem. Unresolved high-consequence signals should affect the diligence limitation, transaction protection or decision timetable.

6. Test schedule logic before accepting completion dates

A schedule is a model of work, dependencies, calendars and constraints. Its finish date is credible only when the underlying network is comprehensive, logically connected, resource-aware and updated with reliable progress. The GAO Schedule Assessment Guide sets out ten practices for a reliable, high-quality schedule and links schedule risk to life-cycle cost.[1]

Automated tests can identify missing predecessors or successors, excessive lags, long activities, hard constraints, negative float, invalid actual dates and broken critical paths. Graph analysis can reveal disconnected clusters and activities whose apparent flexibility depends on incomplete logic. These tests should be calibrated to the planning standard and project phase.

Progress evidence requires separate review. Reported percentage complete can be subjective when the activity lacks a measurable rule. The diligence team should compare progress with quantities, certified milestones, installed equipment, approved drawings, test records, invoices and site evidence. An activity should not become complete merely because its planned finish has passed.

Re-baselining can conceal performance history. The team should compare current and prior baselines, identify moved milestones, added activities, deleted scope, changed calendars and revised logic. The analysis should distinguish approved scope change from recovery assumptions and schedule optimism.

Resource constraints can invalidate a logically sound network. Critical engineering, commissioning or specialist teams may be scheduled across simultaneous activities. Supplier and permit dates may sit outside the native file. The schedule review should connect the network to procurement, staffing, access and interface evidence.

7. Quantify schedule risk as a distribution and causal map

A single forecast date hides uncertainty. Schedule risk analysis should model uncertainty in durations, correlations, discrete events and conditional branches. The output is a distribution of completion outcomes and a map of the drivers that move them.

The model should distinguish uncertainty from management contingency. A duration range can represent ordinary variability. A discrete risk can represent a permit delay, equipment failure during testing or supplier default. A recovery plan should be modelled as an intervention with cost, resource and execution assumptions.

Figure 3. Schedule-risk heat map by workstream and transaction exposure

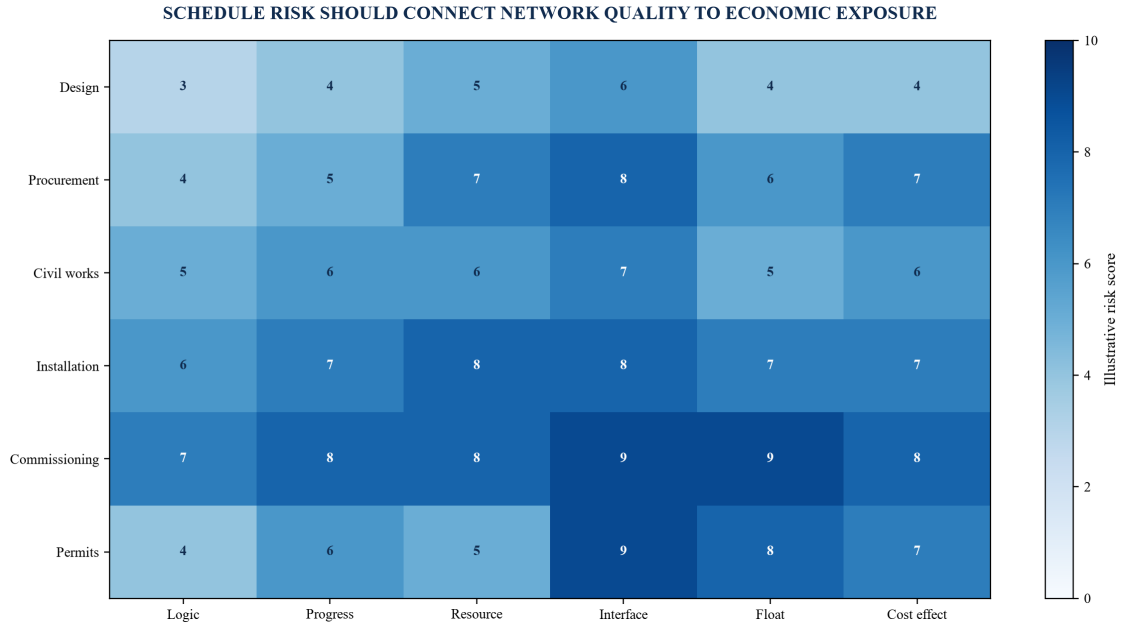

Hypothetical scores illustrate prioritisation only; live assessments require project-specific schedule evidence.

The transaction model should translate time into economics. Delay can extend owner costs, interest during construction, rentals, guarantees and working capital. It can defer revenue, trigger liquidated damages or shorten a contracted operating period. Some impacts are non-linear because a missed seasonal window or permit milestone can cause a larger step change.

Management challenge should focus on the causal chain. If commissioning drives the tail, the team should examine test procedures, system completion, vendor attendance, spares and punch-list closure. If procurement drives the tail, it should examine manufacturing progress, shipping, customs, site readiness and substitute supply. The schedule distribution becomes decision-useful when each driver has evidence and a mitigation owner.

8. Rebuild cost-to-complete from scope, commitments and actuals

The GAO Cost Estimating and Assessment Guide emphasises technical baselines, work breakdown structures, assumptions, data, estimating methods, sensitivity, risk analysis, documentation and updates with actual costs.[2] These practices provide a useful control frame for transaction diligence even when the project is privately financed.

Cost-to-complete should begin with remaining scope, not the seller's forecast total. The team should reconcile budget, actual cost, commitments, accruals, approved changes, pending changes, claims, contingency and forecast. It should identify costs omitted because they sit in another entity, contract or account.

AI can accelerate invoice classification, contract-to-cost mapping and comparison of narrative forecasts with ledger movements. It can flag cost accounts whose remaining forecast is below open commitments, work packages with progress but no accrual, and commitments that lack schedule activities. Each signal requires accounting and engineering review.

Contingency should be analysed by ownership and risk. General contingency, design allowance, escalation, foreign exchange, quantity growth and schedule reserve answer different questions. Drawn contingency should be reconciled to realised risks. Remaining contingency should be

tested against the current risk register and maturity of scope.

The team should distinguish sunk cost from future economic exposure. Historical overspend can indicate weak controls, but the transaction decision depends on remaining cost and recoverability. A seller claim against a contractor may reduce exposure only when entitlement, evidence, counterparty capacity and collection timing support it.

9. Read change orders and claims as a system

Change orders connect scope, schedule, cost, contract rights and cash. Their analysis should extend beyond approved value. The diligence team should map each change to instruction, design revision, notice, quotation, approval, work performed, cost incurred, schedule effect, invoice and collection.

Language models can extract parties, dates, clauses, requested amounts and stated causes from correspondence. The controlled record should retain the original document and distinguish extracted text from legal interpretation. Privilege, confidentiality and jurisdiction-specific advice require explicit handling.

Repeated causes can reveal structural risk. Late design information may drive rework across packages. Access restrictions may create cumulative disruption. A supplier may submit fragmented claims for one underlying event. Clustering can help identify these patterns, followed by engineering and legal review.

The financial model should separate approved, probable, disputed and unsubmitted changes. Revenue recognition, cost accrual, cash collection and legal entitlement may differ. An optimistic claim asset can inflate margin and working capital while creating future funding exposure.

Transaction responses include price adjustment, specific indemnity, escrow, retention, completion support, excluded asset, condition precedent or a revised operating plan. The response should follow the evidence, materiality, recoverability and control of the underlying cause.

10. Connect procurement evidence to schedule and performance

Procurement risk spans specification, supplier capability, manufacturing, quality, logistics, customs, storage, installation and warranty. A purchase order marked placed does not prove that the equipment will arrive, integrate or perform on time.

The data map should connect schedule activities to requisitions, purchase orders, vendor drawings, inspection points, manufacturing status, shipment documents and site receipt. AI-assisted matching can identify missing links and contradictory dates. Physical inspection and supplier confirmation remain necessary for material items.

Long-lead equipment requires evidence of design freeze, approved submittals, production slots, progress payments, quality plans, factory testing and shipping readiness. Supplier financial distress can create hidden exposure even when reported manufacturing progress is high.

Substitution risk matters when a project assumes alternative equipment. The team should test technical equivalence, certification, interface, warranty, lead time and contract approval. A lower-cost substitute can create commissioning delay or lifecycle cost that exceeds the procurement saving.

The transaction documents can allocate identified procurement risks through completion support, specific covenants, retained consideration or closing conditions. The ownership plan should name the person responsible for supplier engagement and interface closure from signing through completion.

11. Underwrite commissioning as an evidence gate

Commissioning converts installed equipment into an operating system. Mechanical completion, energisation, functional testing, integrated testing, performance testing and final acceptance are separate states. Diligence should identify the evidence and authority required for each.

Test results should be compared with design duty, contractual guarantees and actual operating conditions. A passing test at partial load may not support full-load reliability. A short test may not reveal thermal, vibration or control instability. Waivers and punch-list items should be linked to risk and remedy.

AI can compare test sheets, alarm logs and performance curves across equipment populations. It can identify unusual failure sequences or repeated manual overrides. The reviewer should verify sensor calibration, test boundaries, data completeness and the effect of operator intervention.

Interface risk often appears during integrated testing. Equipment that passes independently can fail when connected to power, control, utilities or upstream process conditions. The systems-engineering approach described in NASA's handbook emphasises requirements, interfaces, verification, validation and technical management across the lifecycle.[11]

Closing readiness should distinguish evidence available before signing, evidence expected before completion and evidence that can only arise during ownership. Material post-signing tests may require a condition, holdback, warranty, insurance or funded remediation plan.

12. Build reliability analysis from duty, failure and maintenance evidence

Reliability is the probability that an item performs its required function under stated conditions for a stated period. Diligence should specify the function, operating context, demand, failure definition and observation period before calculating a rate.

Failure history alone can mislead. A low number of failures may reflect limited runtime, hidden standby use, missing work orders or recent replacement. A high number may reflect aggressive detection and good reporting. The analysis should connect failures to exposure, duty, environment, maintenance and configuration.

Figure 4. Reliability curves under alternative evidence conditions

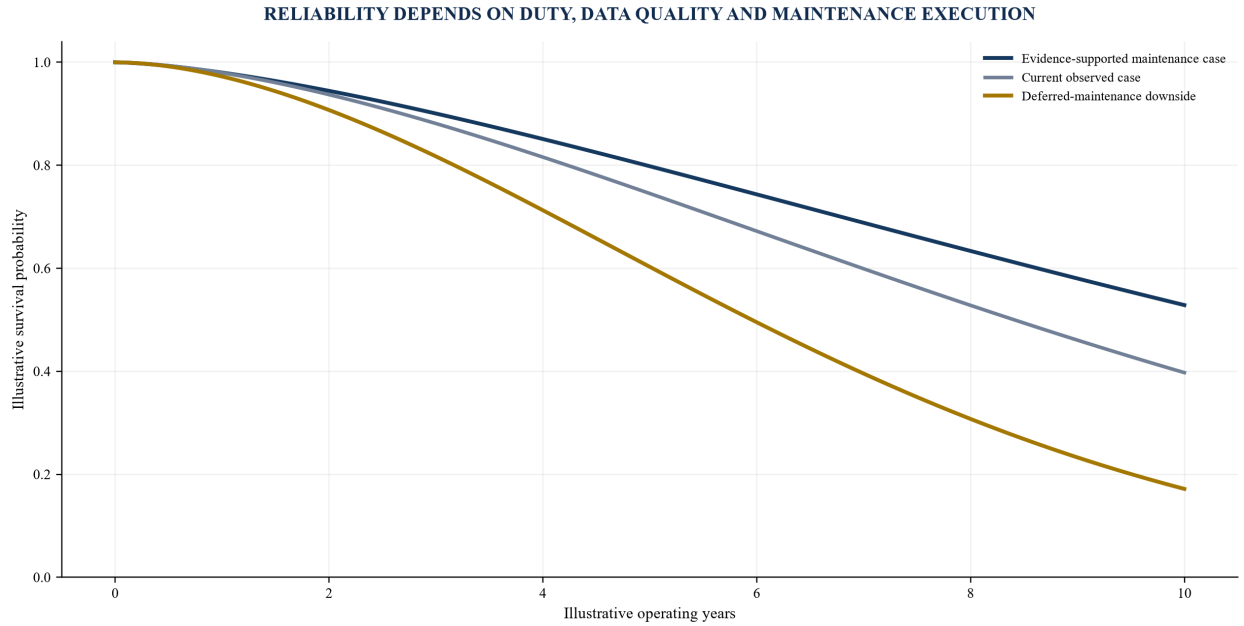

Hypothetical curves demonstrate sensitivity; they are not estimates for an identified asset.

The equipment population should be segmented by design, age, duty and modification state. Pooling unlike assets can hide a weak cohort. Censored observations and replaced items should be handled explicitly. Confidence intervals should be reported where the sample is small.

IEC 61508 addresses functional safety for electrical, electronic and programmable electronic safety-related systems.[10] Safety integrity, reliability and availability are related but distinct. A transaction review should preserve those distinctions and use qualified functional-safety expertise where relevant.

Table 3. Reliability evidence and transaction interpretation

Evidence	Potential signal	Alternative explanation	Diligence response
rising corrective work	deterioration	improved reporting	failure-mode and exposure review
low failure count	strong reliability	low runtime or missing records	normalise exposure and test completeness
repeated alarm reset	unstable operation	faulty sensor	inspect control logic and instrument health
growing spare usage	wear or design issue	stocking-policy change	reconcile issue records and population
deferred inspection	integrity exposure	revised approved interval	verify basis, approval and consequence

Interpretation requires asset-specific engineering judgement and operating context.

13. Convert maintenance history into future cash requirements

Maintenance diligence should reconcile planned work, completed work, overdue work, breakdowns, inspections, defects, spares and contractor capacity. The aim is to estimate the work and cash required to preserve safe and reliable operation under the buyer's plan.

Free-text work orders can be classified into failure modes, causes and actions. The model should retain uncertainty where language is ambiguous. Repeated symptom codes may represent one unresolved root cause. Closure text such as tested okay may not prove that a defect was eliminated.

Backlog should be segmented by safety consequence, production consequence, statutory requirement, asset criticality, age and access window. Monetary estimates should include labour, material, specialist support, outage and consequential work. Backlog value is not simply the sum of work-order estimates because related tasks can overlap or require a common shutdown.

The HSE's asset-integrity work highlights the need for maintenance systems that keep safety-critical elements available and reliable when required.[20] Transaction diligence should test whether inspection and maintenance regimes reflect current hazards, degradation mechanisms and operating duty.

The ownership plan should distinguish catch-up maintenance, recurring maintenance and improvement capital. Catch-up work may affect price or debt capacity. Recurring work belongs in steady-state cash flow. Improvement capital requires a separate business case and should not be used to conceal the cost of restoring basic control.

14. Review operational technology and data continuity

Engineering analytics often relies on historians, controllers, sensors, maintenance platforms and remote vendor access. These systems affect both operation and the evidence used in diligence. A cyber or data-continuity weakness can therefore create safety, production and transaction risk.

NIST SP 800-82 Revision 3 provides guidance for securing operational technology while recognising performance, reliability and safety requirements.[9] The diligence scope should map assets, zones, conduits, remote access, identity, logging, backups, patching, unsupported systems and incident response.

AI models should receive the minimum necessary data through controlled access. Extraction should avoid interfering with production systems. Read-only collection, secure staging, malware scanning, encryption and access logs should be defined. Sensitive process data and intellectual property require contractual and technical controls.

The team should test whether analytical results can be reproduced after completion. Seller-hosted software, proprietary connectors or undocumented tag mappings can create dependency. Transition arrangements should cover source exports, schemas, configuration, credentials, licences and specialist support.

Cyber findings should be linked to operational consequence. A vulnerable remote-access pathway may require closure before signing or completion. A weak asset inventory may require a funded remediation plan. A model should not actively probe production systems unless the scope, safety controls and authorisation expressly permit it.

15. Govern model risk and human review

The AI model register should record purpose, owner, version, data, validation, limitations, access and approved use. It should distinguish deterministic rules, statistical models and generative systems because their failure modes differ.

Validation should test accuracy, completeness, robustness and reproducibility against representative transaction data. It should include known positives, known negatives and difficult edge cases. Performance should be assessed at the unit that matters to the decision, such as activity, work order, equipment tag or finding.

Human review requires competence and independence appropriate to the question. A planning engineer should review schedule logic. A cost engineer should review forecast and contingency. A reliability specialist should review failure analysis. Legal and accounting conclusions should remain with qualified advisers.

The model should expose uncertainty. Language-model extraction can include the source passage and confidence. Anomaly models can show variables contributing to the score. Schedule tests can show the affected activities and paths. Black-box outputs without evidence should remain discovery aids.

The NIST AI RMF critical-infrastructure profile initiative recognises that AI use across information technology, operational technology and industrial control systems requires a repeatable lifecycle approach to trustworthiness.[5] Regulation and contractual obligations depend on jurisdiction and use. The EU AI Act establishes a legal framework for development and use of AI systems within its scope.[16] A cross-border deal should therefore map where the model is developed, operated and used.

16. Translate confirmed findings into valuation and financing

A confirmed engineering finding affects value through cash, timing, risk or optionality. The model should show the causal path. A commissioning delay can extend owner costs and defer revenue. Deferred maintenance can require immediate capital and reduce availability. A reliability weakness can increase spares, outage and insurance cost.

The valuation bridge should avoid double counting. A cost-to-complete increase may already reduce forecast cash flow. The same risk should not also receive a separate multiple discount unless it affects residual uncertainty or future economics. Finance and engineering reviewers should approve the treatment together.

Figure 5. Diligence finding to transaction decision tree

EACH CONFIRMED FINDING SHOULD CHANGE A CONTROLLED TRANSACTION DECISION

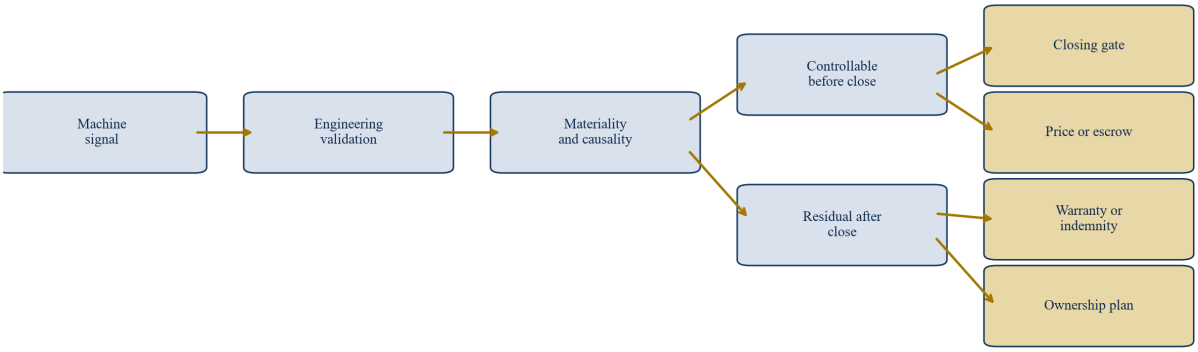

The decision path is illustrative; actual remedies depend on evidence, contract and governing law.

Debt capacity can be affected through remaining capital expenditure, completion risk, availability, maintenance cash, insurance and covenant headroom. The financing model should test the timing of remediation and the party funding it. Seller support has value only when the obligation, amount, duration and credit are credible.

IFRS 3 governs recognition and measurement of assets and liabilities in a business combination and the resulting goodwill or bargain purchase.[13] IFRS 13 provides a framework for fair-value measurement.[15] Engineering findings should be communicated to accounting and valuation specialists with sufficient evidence to assess recognition, measurement and disclosure.

Table 4. Engineering finding to economic and transaction response

Finding	Economic path	Valuation test	Possible transaction response
completion delay	higher cost and deferred revenue	dated cash flow and downside	condition, support or price
cost-to-complete gap	additional funding	sources and uses	price, escrow or commitment
reliability weakness	outage and maintenance cash	availability and lifecycle cost	warranty, reserve or capex plan
OT exposure	disruption and remediation	scenario loss and required spend	closing action or covenant
data limitation	residual uncertainty	sensitivity and decision confidence	protection, further diligence or stop

Treatment depends on materiality, control, accounting, contract and recovery evidence.

17. Draft transaction protections from causal evidence

Transaction protection should follow the identified risk and the party best able to control or absorb it. Generic warranties can be weak when the finding concerns a specific project, asset or data limitation. The diligence record should support precise drafting by qualified legal advisers.

A closing condition can require completion of a test, delivery of evidence, consent, permit or remediation. A price adjustment can recognise a measurable funding gap. An escrow or retention can secure a defined exposure. A warranty or indemnity can allocate risk that survives completion. A covenant can govern work after signing.

Remedies require triggers and measurement. If protection depends on reliability, the document should define the equipment, function, period, operating conditions, failure event, evidence and calculation. If it depends on project completion, it should define milestones, acceptance, exclusions and long-stop consequences.

IAS 37 addresses provisions, contingent liabilities and contingent assets.[14] Legal rights and accounting recognition are separate questions. The transaction team should preserve that distinction when converting an engineering finding into a model or disclosure.

The evidence package should be usable in negotiation. It should contain the confirmed finding, source documents, analysis, management response, specialist conclusion, economic effect and proposed action. The model output alone is insufficient.

18. Build the ownership plan before signing

Some engineering risks can only be resolved during ownership. The buyer should convert each accepted risk into an initiative with an owner, scope, budget, dependency, milestone, evidence and decision right before signing.

The first hundred days may include schedule re-baselining, cost-to-complete validation, critical maintenance, spares procurement, OT access control, configuration cleanup and supplier engagement. Work should be sequenced around safety, continuity and cash rather than a generic integration calendar.

Data continuity is a Day-One requirement. The buyer should secure access to native schedule files, cost histories, asset registers, maintenance records, historian exports, model documentation and configuration. Transition services should include response times, security and exit arrangements.

Management capability is part of the engineering risk. The team should identify planners, cost engineers, reliability specialists, control engineers and maintenance leaders whose knowledge is required. Retention and succession actions should follow from verified role dependence.

The investment committee should see the total ownership cost. Purchase price, completion funding, catch-up maintenance, systems remediation and contingency should appear in one sources-and-uses and value bridge. An attractive headline multiple can conceal a large post-completion restoration programme.

19. Govern diligence through evidence gates and decision logs

The diligence office should operate a weekly evidence cadence. Each material question has an owner, required evidence, status, conclusion, economic effect and proposed decision. Open items should be ranked by materiality and time to resolve.

The model team should publish controlled releases. Each release records source snapshots, code or model version, parameters, tests and changed outputs. Material changes should be explained to

engineering and transaction leads before they enter a board paper.

Table 5. Governance gates for AI-assisted engineering diligence

Gate	Required evidence	Decision owner	Failure response
Scope	thesis, perimeter and question register	deal lead	revise mandate or defer test
Data	provenance, access and quality map	engineering lead	qualify analysis or obtain evidence
Model	validation, version and permitted use	model owner	restrict use or remediate
Finding	causality, materiality and reviewer sign-off	specialist lead	reject, defer or escalate
Deal	economic bridge and protection	investment committee	reprice, restructure or stop

The gate design should be calibrated to transaction timetable and delegated authority.

The decision log should record why a finding changed or did not change the transaction. This protects against late-stage memory loss and helps the ownership team understand accepted risks. It also creates feedback for future transactions.

Conflicts should be visible. A seller may control data extraction. A vendor may benefit from identifying replacement work. A model provider may promote its own accuracy. Reviewer independence, access and compensation should be assessed where material.

The board pack should separate confirmed findings, unresolved limitations and management scenarios. It should state which conclusions rely on physical inspection, third-party confirmation or future testing. This structure supports a decision under uncertainty without disguising missing evidence.

20. Create a repeatable engineering-diligence control system

AI-assisted engineering diligence becomes repeatable when the organisation preserves a controlled method rather than a collection of prompts and spreadsheets. The method connects scope, data, analysis, review, economics, transaction protection and ownership action.

Reusable components include data dictionaries, identifier crosswalks, schedule tests, cost reconciliations, maintenance taxonomies, reliability notebooks, evidence templates and decision logs. They should be adapted to asset type and transaction thesis. Standardisation should preserve engineering judgement and local regulatory requirements.

The system should learn from outcomes. After completion, actual cost, schedule, reliability and remediation results can be compared with diligence findings. This feedback can reveal optimistic thresholds, weak data sources, missed failure modes and valuable early indicators. Model improvement should use controlled and legally permitted data.

Portfolio learning should preserve transaction context. A schedule signal from a greenfield project may not transfer directly to a brownfield shutdown. Failure modes in rotating equipment may have little relevance to software-controlled logistics assets. The library should therefore retain asset class, lifecycle phase, operating duty, data completeness, intervention and outcome alongside each finding. These attributes help reviewers select comparable evidence and avoid

inappropriate generalisation.

Quality assurance should cover the complete analytical chain. Source-population checks test whether all required records were received. Transformation checks test identity, units, dates and versions. Model checks test expected behaviour and edge cases. Finding checks test source support, reviewer approval and economic translation. Release checks test that the board pack, valuation model and transaction documents use the same approved finding set. Exceptions should remain visible until resolved or accepted by the authorised decision maker.

The operating model also needs a secure retention policy. Some evidence may be required for transaction negotiation, financing, insurance, warranty administration or future disputes. Other target data may need deletion after a defined period. The diligence office should record retention, access, legal hold and disposal rules by data class. Model-training use should require separate authority and should respect contractual, privacy, competition and intellectual-property restrictions.

The ultimate measure is decision quality. The analysis should surface material risks early enough to change price, structure, protection, readiness or ownership execution. It should also dismiss immaterial noise efficiently. Coverage, traceability, review time, finding conversion and post-deal accuracy provide a balanced performance set.

For the buyer, the output is a defensible bridge from engineering evidence to value and risk. For lenders, it is a clearer view of completion, maintenance and cash exposure. For management, it is a prioritised ownership plan. For the board, it is an auditable basis for deciding whether and how to sign.

References

[1] U.S. Government Accountability Office, Schedule Assessment Guide: Best Practices for Project Schedules, GAO-16-89G, 2015, https://www.gao.gov/products/gao-16-89g

[2] U.S. Government Accountability Office, Cost Estimating and Assessment Guide: Best Practices for Developing and Managing Program Costs, GAO-20-195G, 2020, https://www.gao.gov/products/gao-20-195g

[3] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, NIST AI 100-1, 2023, https://doi.org/10.6028/NIST.AI.100-1

[4] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, 2024, https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

[5] National Institute of Standards and Technology, Concept Note: Trustworthy Use of AI in Critical Infrastructure Profile, 2026, https://www.nist.gov/system/files/documents/2026/04/08/Draft%20Concept%20Note_%20Development%20of%20the%20NIST%20AI%20RMF%20Trustworthy%20Use%20of%20AI%20in%20Critical%20Infrastructure%20Profile.pdf

[6] International Organization for Standardization, ISO/IEC 42001:2023 Artificial Intelligence Management System, https://www.iso.org/standard/42001

[7] International Organization for Standardization, ISO 55001:2024 Asset Management System Requirements, https://committee.iso.org/sites/tc251/home/projects/published/iso-55001.html

[8] International Organization for Standardization, ISO 55013:2024 Guidance on the Management of Data Assets, https://www.iso.org/management-system-standards-list.html

[9] National Institute of Standards and Technology, Guide to Operational Technology Security, NIST SP 800-82 Revision 3, 2023, https://doi.org/10.6028/NIST.SP.800-82r3

[10] International Electrotechnical Commission, IEC 61508-2:2010 Functional Safety Requirements for Electrical, Electronic and Programmable Electronic Safety-Related Systems, https://webstore.iec.ch/en/publication/5516

[11] National Aeronautics and Space Administration, NASA Systems Engineering Handbook, NASA/SP-2016-6105 Revision 2, https://www.nasa.gov/reference/systems-engineering-handbook/

[12] National Aeronautics and Space Administration, Expanded Guidance for NASA Systems Engineering, Volume 1, NASA/SP-2016-6105/SUPPL, 2016, https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20170007238.pdf

[13] IFRS Foundation, IFRS 3 Business Combinations, https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/

[14] 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/

[15] IFRS Foundation, IFRS 13 Fair Value Measurement, https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/

[16] European Union, Regulation (EU) 2024/1689, Artificial Intelligence Act, https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[17] Organisation for Economic Co-operation and Development, OECD AI Principles, https://oecd.ai/en/ai-principles

[18] Cybersecurity and Infrastructure Security Agency, Securing Industrial Control Systems, https://www.cisa.gov/sites/default/files/publications/Securing_Industrial_Control_Systems_S508C.pdf

[19] National Aeronautics and Space Administration, MSFC Systems Engineering Handbook, MSFC-HDBK-3173, https://standards.nasa.gov/standard/msfc/msfc-hdbk-3173

[20] UK Health and Safety Executive, Key Programme 3: Asset Integrity Programme, https://www.hse.gov.uk/offshore/assets/docs/kp3.pdf

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.