

Technical Diligence for AI Targets: Models, Data, Compute, Security and Human Review

A Control-Based Framework for Capability, Dependencies, Cost and Acquisition Risk

Authored by Chennakeshav Adya, Independent Researcher

September 2026

Abstract

An acquisition target can describe itself as an artificial-intelligence company while its commercial product depends on rented models, poorly governed data, scarce compute, fragile integrations and manual intervention. This paper develops a control-based technical-diligence framework for buyers, boards, investment committees and lenders. It maps the complete AI system perimeter across models, applications, agents, data pipelines, infrastructure, interfaces and human decision points. Product claims are reconciled to production telemetry, customer use and contractual commitments. Model inventory, provenance, licences, version history and reproducible evaluations establish what the target owns and what its systems can reliably do. Robustness testing covers distribution shift, adversarial inputs, hallucination, prompt injection, tool misuse and unsafe outputs. The data review traces training, fine-tuning, retrieval, inference, feedback and evaluation datasets, testing rights, quality, leakage, bias, privacy and deletion controls. Compute diligence quantifies accelerator capacity, cloud concentration, utilisation, cost, regional resilience, egress and portability. Software supply-chain, identity, application, infrastructure and agentic-workflow controls reveal security exposure. Human oversight is tested through authority, competence, override, escalation and appeal evidence. Monitoring, incident history, governance, testing independence, regulatory classification, transparency and intellectual-property rights complete the control picture. Third-party model concentration, customer change rights, key-person dependency and engineering velocity determine continuity. The analysis separates proprietary assets from replaceable capability and rebuilds unit economics for training, inference, retrieval, observability and review. Remediation cost and commercial stress cases flow into revenue, margin, capital needs, valuation and transaction protections. Day-One controls and a one-hundred-day roadmap allocate owners, budgets, dependencies and decision gates. Five figures, five tables, eight frequently asked questions and twenty-six primary or authoritative sources support implementation. Numerical scores are illustrative analytical examples. Every conclusion depends on verified technical, commercial, contractual, regulatory, data, security, financial, tax and transaction facts and requires authorised professional advice.

JEL Classification: G32, G34, L86, O32, O33

Keywords: artificial intelligence, technical diligence, model risk, data rights, compute economics, cybersecurity, human oversight, M&A

1. Define the acquisition question

Translate the investment thesis into testable claims about capability, defensibility, cost, safety and control.

The AI diligence team should reconcile investment case, product roadmap, commercial model, architecture and risk appetite. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI diligence mandate.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

2. Map the AI system perimeter

Identify models, applications, agents, data pipelines, infrastructure, interfaces and human decision points.

The AI diligence team should reconcile architecture diagrams, repositories, inventories, deployment records and contracts. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI system map.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

3. Reconcile product claims

Tie marketed features and demonstrations to production behaviour, customer use and contractual commitments.

The AI diligence team should reconcile product collateral, demos, telemetry, contracts, support records and customer evidence. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a claims-to-evidence register.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

4. Classify roles and obligations

Determine provider, deployer, importer, distributor and downstream responsibilities by jurisdiction and use case.

The AI diligence team should reconcile legal-entity map, product flows, customer terms, regulatory analysis and policies. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an obligations matrix.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

5. Build the model inventory

Record foundation, fine-tuned, classical and third-party models, versions, owners, purposes and deployment status.

The AI diligence team should reconcile model registry, repositories, endpoints, bills of materials and change logs. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a governed model register.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

6. Establish model provenance

Trace weights, training method, licences, source lineage, modifications and approval history.

The AI diligence team should reconcile licences, model cards, commits, training records, vendor terms and attestations. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a model-provenance chain.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

7. Test model performance

Reproduce material accuracy, reliability, latency, calibration and task-specific evaluation claims.

The AI diligence team should reconcile evaluation suites, test sets, production telemetry, baselines and release gates. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a reproducible performance dossier.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

8. Test robustness and failure modes

Probe distribution shift, adversarial inputs, hallucination, prompt injection, tool misuse and unsafe outputs.

The AI diligence team should reconcile red-team results, incident logs, adversarial tests, safeguards and monitoring. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a model-risk assessment.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.


Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

Table 1. AI system perimeter

Layer	Primary evidence	Core risk
model	registry and provenance	unowned capability
data	lineage and rights	unlawful or weak inputs
compute	contracts and telemetry	cost and concentration
workflow	permissions and review	unsafe execution

Illustrative programme design; company-specific facts and authorised advice govern.

Figure 1. System evidence confidence



Values are illustrative readiness indices and require company-specific evidence.

9. Map the data estate

Identify training, fine-tuning, retrieval, inference, feedback and evaluation datasets and their owners.

The AI diligence team should reconcile data catalogues, lineage, schemas, access logs, contracts and repositories. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI data map.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

10. Verify data rights

Test collection basis, consent, contractual permission, copyright position, retention and geographic restrictions.

The AI diligence team should reconcile licences, privacy notices, consents, customer terms, DPIAs and legal advice. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a data-rights schedule.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

11. Measure data quality

Assess completeness, representativeness, labelling, contamination, drift, leakage and bias.

The AI diligence team should reconcile dataset statistics, labelling protocols, quality tests, benchmarks and exceptions. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a data-quality report.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

12. Control personal and sensitive data

Trace minimisation, purpose limitation, access, deletion, rights handling and automated-decision safeguards.

The AI diligence team should reconcile records of processing, DPIAs, access controls, deletion tests and request logs. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a privacy-control matrix.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

13. Map compute dependencies

Quantify accelerators, cloud services, regions, capacity reservations, networking, storage and portability.

The AI diligence team should reconcile cloud bills, contracts, capacity plans, utilisation data and architecture. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a compute-dependency map.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

14. Rebuild unit economics

Calculate training, fine-tuning, inference, retrieval, observability and human-review cost by product and customer.

The AI diligence team should reconcile usage telemetry, vendor invoices, price books, workload data and allocation rules. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI cost-to-serve model.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

15. Stress compute availability

Model accelerator scarcity, regional failure, vendor limits, price changes, egress and migration constraints.

The AI diligence team should reconcile capacity contracts, quotas, resilience tests, price scenarios and exit plans. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a compute-resilience stress test.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

16. Test software supply chain

Inventory code, packages, containers, APIs, models and data dependencies with licences and vulnerabilities.

The AI diligence team should reconcile SBOMs, model bills of materials, scans, repositories and vendor registers. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI supply-chain dossier.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.


Table 2. Model and data tests

Dimension	Required test	Decision use
performance	reproducible evaluation	product claims
robustness	adversarial and shift testing	downside
rights	licence and purpose review	operability
quality	drift and leakage tests	reliability

Illustrative programme design; company-specific facts and authorised advice govern.

Figure 2. Model and data readiness

FIGURE 2. MODEL AND DATA READINESS



Values are illustrative readiness indices and require company-specific evidence.

17. Test identity and access

Review privileged access, service accounts, model endpoints, agent permissions, secrets and segregation.

The AI diligence team should reconcile IAM configuration, access reviews, key vaults, logs and penetration tests. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI access-control assessment.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

18. Test application security

Probe prompt injection, insecure output handling, data exfiltration, excessive agency and denial of service.

The AI diligence team should reconcile threat models, secure design reviews, tests, incidents and compensating controls. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI application-security report.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

19. Test infrastructure security

Assess cloud configuration, network boundaries, encryption, backups, resilience and incident response.

The AI diligence team should reconcile cloud posture, architecture, key management, recovery tests and playbooks. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an infrastructure-control report.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

20. Review agentic workflows

Trace tool calls, permissions, memory, escalation, transaction limits and kill switches.

The AI diligence team should reconcile agent specifications, tool inventories, traces, approval rules and failure tests. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an agent-control map.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

21. Assess human review

Determine where judgement, override, appeal, escalation and accountable approval are necessary and effective.

The AI diligence team should reconcile workflow maps, sampling, reviewer training, overrides, appeals and outcomes. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a human-oversight assessment.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

22. Test monitoring and observability

Verify logging, drift detection, safety metrics, alerts, lineage, retention and investigation capability.

The AI diligence team should reconcile telemetry, dashboards, alerts, log samples, incident tickets and runbooks. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI observability dossier.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

23. Reconstruct the incident record

Identify model, data, security, privacy, safety and availability events and management response.

The AI diligence team should reconcile incident registers, customer notices, root-cause analyses, remediation and claims. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI incident history.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

24. Assess governance

Test ownership, risk classification, policies, committees, release authority and independent challenge.

The AI diligence team should reconcile governance charters, policies, minutes, risk registers and approvals. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI governance assessment.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.


Table 3. Control architecture

Control	Evidence	Failure signal
identity	privileged-access review	orphaned credentials
security	threat tests	exfiltration path
oversight	override samples	rubber-stamp review
incidents	closed-loop remediation	repeat event

Illustrative programme design; company-specific facts and authorised advice govern.

Figure 3. Security and control readiness

FIGURE 3. SECURITY AND CONTROL READINESS



Values are illustrative readiness indices and require company-specific evidence.

25. Review testing governance

Evaluate independence, test coverage, benchmark relevance, version control and acceptance thresholds.

The AI diligence team should reconcile test plans, evaluation protocols, sign-offs, exceptions and release evidence. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a testing-control framework.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

26. Map regulatory exposure

Assess current and scheduled obligations across product, sector and geography.

The AI diligence team should reconcile regulatory inventory, classifications, legal analysis, technical files and roadmaps. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a regulatory exposure map.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

27. Review transparency controls

Test notices, disclosures, content provenance, synthetic-media labelling and customer documentation.

The AI diligence team should reconcile user interfaces, model cards, content credentials, terms and audit samples. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a transparency-control report.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

28. Assess intellectual property

Test ownership and licence rights in code, models, data, inventions, brands and generated outputs.

The AI diligence team should reconcile employment terms, assignments, licences, patents, claims and disputes. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI intellectual-property schedule.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

29. Assess third-party concentration

Measure dependence on foundation models, clouds, data vendors, open-source maintainers and key contractors.

The AI diligence team should reconcile vendor register, spend, contracts, architecture, substitution tests and exit plans. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a vendor-concentration analysis.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

30. Test customer dependency and change rights

Map model substitutions, performance promises, data uses, audit rights, indemnities and termination triggers.

The AI diligence team should reconcile customer contracts, amendments, security schedules, correspondence and claims. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI contract-risk matrix.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

31. Assess technical team dependency

Identify critical knowledge, key-person concentration, documentation gaps and retention risk.

The AI diligence team should reconcile organisation charts, repositories, ownership maps, interviews and succession plans. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a technical-capability map.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

32. Review development velocity

Reconcile roadmap delivery, release frequency, defect escape, rework and platform debt.

The AI diligence team should reconcile roadmaps, commits, releases, tickets, incidents and engineering metrics. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a delivery-capability analysis.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.


Table 4. Defensibility tests

Claim	Diligence question	Evidence
proprietary model	what is owned	weights and rights
unique data	can it be used	lineage and licences
cost advantage	does it persist	unit economics
workflow moat	what would switching require	customer evidence

Illustrative programme design; company-specific facts and authorised advice govern.

Figure 4. Defensibility confidence

FIGURE 4. DEFENSIBILITY CONFIDENCE



Values are illustrative readiness indices and require company-specific evidence.

33. Separate durable assets from rented capability

Distinguish proprietary data, workflow integration, evaluation, distribution and know-how from replaceable APIs.

The AI diligence team should reconcile architecture, contracts, customer evidence, benchmarks and switching tests. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is a defensibility bridge.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

34. Quantify remediation cost

Cost mandatory fixes, documentation, security uplift, data remediation, re-platforming and governance.

The AI diligence team should reconcile gap register, estimates, vendor quotes, resourcing plans and dependencies. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI remediation budget.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

35. Stress revenue and margin

Model model-price changes, compute inflation, quality degradation, outages, compliance delay and customer churn.

The AI diligence team should reconcile unit economics, contracts, pipeline, scenarios, capacity and risk findings. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI commercial stress test.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

36. Translate findings into valuation

Adjust forecasts, margins, capital needs, terminal assumptions and multiples for technical evidence and risk.

The AI diligence team should reconcile valuation model, diligence findings, scenarios, remediation and market evidence. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI value bridge.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

37. Design transaction protections

Allocate identified risks through price, holdback, escrow, indemnity, warranty, covenant or condition.

The AI diligence team should reconcile risk register, valuation, legal drafting, insurance and negotiation. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI protection matrix.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

38. Design Day-One controls

Protect credentials, repositories, model endpoints, vendors, data access and incident authority at close.

The AI diligence team should reconcile access plan, TSA, integration design, approvals and response playbooks. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI Day-One control plan.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

39. Build the one-hundred-day roadmap

Sequence evidence closure, security, governance, platform, data, people and commercial integration.

The AI diligence team should reconcile gap register, owners, budget, dependencies, milestones and board reporting. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI value-and-control roadmap.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.

40. Issue the diligence conclusion

State verified capability, dependencies, cost, compliance, downside, remedies and decision conditions.

The AI diligence team should reconcile reconciled evidence, tests, expert advice, approvals and transaction documents. Each conclusion records the accountable owner, source, system version, use case, evidence, financial consequence, control and unresolved exception. The immediate output is an AI technical diligence certificate.

AI capability and operability must be proved through reconciled technical, commercial and control evidence. Reviewers test performance, provenance, rights, security, cost, human oversight and downside response against native records and observed behaviour. System versions, use cases, transaction perimeter, contractual rights and applicable law control each conclusion.

Material gaps require an owner, corrective action, validation test, budget, advice and decision date. Consequences flow through product continuity, revenue, margin, cash, capital needs, valuation, compliance and transaction protection. Residual risk remains visible until evidence is reconciled, remedies are executable and authorised decision-makers approve the next gate.


Table 5. Acquisition decision certificate

Dimension	Required conclusion	Evidence
capability	claims reproduced	controlled tests
operability	rights and dependencies mapped	registers and contracts
economics	cost and downside modelled	telemetry and scenarios
control	remedies funded and owned	transaction and operating plan

Illustrative programme design; company-specific facts and authorised advice govern.

Figure 5. Acquisition readiness

FIGURE 5. ACQUISITION READINESS



Values are illustrative readiness indices and require company-specific evidence.

References

- [1] National Institute of Standards and Technology, AI Risk Management Framework, <https://www.nist.gov/itl/ai-risk-management-framework>
- [2] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, <https://doi.org/10.6028/NIST.AI.600-1>
- [3] National Institute of Standards and Technology, AI RMF Playbook, <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>
- [4] National Institute of Standards and Technology, AI Resource Center, <https://airc.nist.gov/>
- [5] National Institute of Standards and Technology, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, <https://doi.org/10.6028/NIST.AI.100-2e2023>
- [6] National Institute of Standards and Technology, Cybersecurity Framework 2.0, <https://doi.org/10.6028/NIST.CSWP.29>
- [7] National Institute of Standards and Technology, Secure Software Development Framework, <https://doi.org/10.6028/NIST.SP.800-218>
- [8] European Union, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence, <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [9] European Commission, AI Act regulatory framework and implementation timeline, <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [10] European Commission, Guidelines for providers of general-purpose AI models, <https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers>
- [11] European Commission, General-Purpose AI Code of Practice, <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>
- [12] European Union, General Data Protection Regulation, <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [13] United Kingdom Information Commissioner's Office, Guidance on AI and data protection, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>

- [14] United Kingdom Information Commissioner's Office, AI and data protection risk toolkit, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/ai-and-data-protection-risk-toolkit/>
- [15] UK National Cyber Security Centre, Guidelines for secure AI system development, <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
- [16] Cybersecurity and Infrastructure Security Agency, Secure by Design, <https://www.cisa.gov/securebydesign>
- [17] MITRE, ATLAS Adversarial Threat Landscape for Artificial-Intelligence Systems, <https://atlas.mitre.org/>
- [18] OWASP Foundation, OWASP Top 10 for Large Language Model Applications, <https://owasp.org/www-project-top-10-for-large-language-model-applications/>
- [19] OECD, Recommendation of the Council on Artificial Intelligence, <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- [20] OECD, AI Principles, <https://oecd.ai/en/ai-principles>
- [21] International Organization for Standardization, ISO/IEC 42001 Artificial intelligence management systems, <https://www.iso.org/standard/81230.html>
- [22] International Organization for Standardization, ISO/IEC 23894 Artificial intelligence risk management, <https://www.iso.org/standard/77304.html>
- [23] IFRS Foundation, IFRS 3 Business Combinations, <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
- [24] IFRS Foundation, IAS 38 Intangible Assets, <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
- [25] IFRS Foundation, IAS 36 Impairment of Assets, <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
- [26] International Valuation Standards Council, IVS 200 Businesses and Business Interests, <https://www.ivsc.org/standards/>

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.