

AI Due Diligence Checklist

90 buy-side checks for diligencing an AI company, grouped by the five layers of the AI Target Audit — Data, Model, Infrastructure, Output and Governance — plus the three cross-cutting groups the audit feeds: customer Contracts, the research Team, and the inference Economics. Every line names the evidence to request, the reviewer who reads it, its disclosure wave, and whether the finding walks the deal, moves price, or simply confirms posture. Free, no signup, and every line traces to [the AI due diligence guide](#).

AI Due Diligence Checklist (2026)

Version 2026-09-04 · 90 checks · <https://www.peony.ink/tools/ai-due-diligence-checklist>

Data Layer (0 of 13 received)

Training-data provenance, license trail, PII exposure and scrape risk — the heaviest of the five layers and the long pole at roughly two to four weeks, because per-dataset traceability cannot be parallelized past the target's own engineers.

- ☐ **Request a per-dataset training-data manifest: source, acquisition date, license terms, intended-use scope, licensee-of-record.**

Provenance is the Data Layer's long pole; a manifest that has to be reconstructed from engineering memory is the single biggest cost swing in the whole review.

Evidence: Per-dataset training-data manifest

Wave 3 — confirmatory

Data-protection counsel

IP counsel

Price impact

- ☐ **Reject any manifest line reading 'publicly available web data' with no further documentation, and trace it back to a named source instead.**

'Publicly available web data' is the canonical unacceptable evidence line: it hides both scrape risk and license-mismatch behind a single manifest entry.

Evidence: Source-of-truth identifier per dataset

Wave 3 — confirmatory

Data-protection counsel

IP counsel

Price impact

- ☐ **Trace every dataset to an acquisition receipt and purchase trail, not just to a license — the two are different documents.**

Bartz turned on how the copies were obtained rather than on what the license said afterwards, so the receipt is the exposure-defining artifact.

Evidence: Acquisition receipts and data-acquisition purchase trail

Wave 3 — confirmatory

IP counsel

Data-protection counsel

Price impact

- ☐ **Ask the binary question: did any training corpus include Library Genesis, Anna's Archive, Pirate Library Mirror or Z-Library content?**

A documented yes scores 1 (Red) on the Data Layer regardless of how clean the other four tests are, and it is the most quantifiable downside in AI diligence.

Evidence: Shadow-library screening record plus engineering-team interview notes

Wave 3 — confirmatory

IP counsel

Data-protection counsel

Walk risk

- ☐ **If shadow-library content is confirmed, test the cure: settlement payment, retraining from clean corpora, and customer notification.**

A partial cure is no cure — without settlement payment, retraining from clean corpora and customer notification the Red score stands and the deal moves to asset-only or walk territory.

Evidence: Post-detection remediation plan, retraining log, customer notification record

Wave 3 — confirmatory

IP counsel

Deal team

Walk risk

- ☐ **Multiply the Bartz per-work baseline — about \$3,000 per book across roughly 500,000 works — against the likely-pirated corpus volume.**

It is the only per-work damages number a court has actually approved, so it converts a qualitative scrape finding into a sized indemnity ask.

Evidence: Corpus volume count for any dataset with shadow-library exposure

Wave 3 — confirmatory

IP counsel

Deal team

Price impact

source

- ☐ **Check each dataset license was effective on the training-run date, and flag any license later revoked or downgraded.**

Training completed under a license that no longer exists is a live IP exposure the seller usually has not priced.

Evidence: Dataset licenses with effective dates, matched to training-run dates

Wave 3 — confirmatory

IP counsel

Price impact

- ☐ **Check for attribution-gap: does the product owe attribution to a dataset source that was never credited?**

Attribution and derivative-works clauses are the quiet half of dataset licensing and are cheap to cure before close, expensive after.

Evidence: Dataset cards and license attribution requirements

Wave 3 — confirmatory

IP counsel

Confirmatory

- ☐ **Require a PII inventory across every training dataset, plus the DPIA for any dataset containing personal data.**

PII in training corpora creates GDPR Article 17, CCPA and state biometric-privacy exposure; no inventory scores the layer 1-2 on its own.

Evidence: PII inventory and data-protection impact assessment (DPIA)

Wave 3 — confirmatory

Data-protection counsel

Price impact

- ☐ **Walk the deletion-on-request workflow from production data through the training corpora to every fine-tune derived from them.**

The chain usually breaks at the fine-tune, which is where a deletion promise quietly stops being true.

Evidence: Deletion-on-request workflow documentation and deletion log

Wave 3 — confirmatory

Data-protection counsel

Price impact

- ☐ **Ask for documented evidence of the retraining or unlearning approach used to honor deletions — a corpus deletion is not a model deletion.**

It is the technical reality behind every erasure commitment, and the absence of a documented approach is a scored finding.

Evidence: Model-retraining or unlearning procedure with execution records

Wave 3 — confirmatory

Data-protection counsel

AI/ML advisor

Price impact

- ☐ **Collect opt-out compliance records: robots.txt AI-crawler directives honored, DMCA notices received, and GDPR Article 17 requests received.**

Opt-out posture is diagnostic test five and shows whether the target responds to rightsholders on a process or on an ad-hoc basis.

Evidence: Opt-out, DMCA and GDPR Article 17 compliance logs

Wave 3 — confirmatory

Data-protection counsel

IP counsel

Confirmatory

- ☐ **Review the user-log preservation policy against litigation-hold exposure, not just the internal retention schedule.**

The January 2026 NYT v. OpenAI order compelling a 20-million-log sample made preservation an obligation that flows from court orders, with real infrastructure cost.

Evidence: User-log retention and litigation-preservation policy

Wave 3 — confirmatory

Data-protection counsel

Deal team

Confirmatory

Model Layer (0 of 12 received)

Proprietary versus licensed models, license-cliff terms, fine-tune lineage and the model-risk register — the license-cliff test is the most material because the consequence is binary: either the product is in license compliance or it is not.

- ☐ **Require a model inventory covering every foundation model, fine-tune, embedded open-weights model, API model and internal ops model.**

The inventory is the LOI-wave artifact the rest of the Model Layer hangs off; a team that can name only three models out of ten scores 1-2 and is hiding a cliff.

Evidence: Model inventory with name, license, version, training-data summary, intended use and in-production flag

Wave 2 — LOI

AI/ML advisor

IP counsel

Price impact

- ☐ **Ask which Llama versions and weights are in the stack, whether they are fine-tuned or merged, and whether Llama sits in the inference path.**

It scopes the whole license-cliff question — a Llama model outside the inference path is a different finding from one inside it.

Evidence: Base-model licenses and deployment map by product surface

Wave 2 — LOI

AI/ML advisor

IP counsel

Price impact

[source](#)

- ☐ **Test the 700-million-monthly-active-user threshold at section 2 of the Meta Llama Community License, measured on the version release date.**

Above the threshold a licensee must request a separate license Meta may grant at its sole discretion, so the consequence is binary compliance or violation.

Evidence: MAU measurement methodology and the applicable Llama license text

Wave 3 — confirmatory

IP counsel

AI/ML advisor

Walk risk

[source](#)

- ☐ **Re-run the 700M MAU test against the acquirer's combined user base — the clause reaches Licensee or Licensee's affiliates.**

A sub-scale target inside a large buyer can lose the right to adopt the next Llama release even though it was comfortably compliant standalone.

Evidence: Post-close combined MAU analysis prepared by counsel

Wave 3 — confirmatory

IP counsel

Deal team

Price impact

- ☐ **Check the Llama 3.1 and Llama 4 naming condition at sections 1.b.i and 1.b.iii on any model improved with Llama Materials or their outputs.**

The naming and attribution conditions bind regardless of MAU, so a compliant-on-scale target can still be non-compliant on its product names.

Evidence: Model naming register cross-checked against the governing license version

Wave 3 — confirmatory

IP counsel

Price impact

[source](#)

- ☐ **Identify which Llama license version governs each fine-tune run before judging cross-model training.**

Llama 2 and 3 prohibited improving other LLMs; 3.1 and 4 dropped that and replaced it with the naming condition, so most memos get this backwards.

Evidence: Fine-tune run log matched to the license version in force on the run date

Wave 3 — confirmatory

IP counsel

AI/ML advisor

Price impact

- ☐ **Write open-weights under a community license in the memo, never 'Llama is open source'.**

Meta's instrument carries a use-based commercial trigger and naming conditions, and that difference is the whole finding.

Evidence: Memo language review against the license text

Wave 3 — confirmatory

IP counsel

Confirmatory

- ☐ **Run the parallel license review on Mistral commercial terms and on the Anthropic AUP for any enterprise customer in regulated use cases.**

Llama is the loudest cliff but not the only one; the same review has to cover every third-party model in the stack.

Evidence: Third-party model license and acceptable-use policy file

Wave 3 — confirmatory

IP counsel

Confirmatory

- ☐ **Require fine-tune lineage for every fine-tuned model: base model, fine-tune data, license on the output, and derivative-work restrictions.**

Lineage is what proves each fine-tune is training-data-traceable to a license-compliant source; an engineering admission of cross-model training scores 1-2.

Evidence: Fine-tune lineage documentation

Wave 3 — confirmatory

AI/ML advisor

IP counsel

Price impact

- ☐ **Request the model-card archive: intended use, performance, bias and safety evaluation, training-data summary, footprint and limitations.**

Cards for every production model score 5, flagship-only scores 3, none scores 1-2 — and GPAI obligations have required the underlying documentation since 2 August 2025.

Evidence: Model-card archive for every production model

Wave 3 — confirmatory

AI/ML advisor

Price impact

- ☐ **Request eval-suite results with the held-out test design, not just headline benchmark scores.**

The held-out design is what tells you whether the evaluation infrastructure is a real asset or a screenshot.

Evidence: Eval-suite results and held-out test design

Wave 3 — confirmatory

AI/ML advisor

Confirmatory

- ☐ **Require a model-risk-management register — OCC 2011-12-aligned for regulated targets, NIST-aligned otherwise — reviewed on a stated cadence.**

Quarterly review scores 5, a static one-time register scores 3, none scores 1-2, and regulated buyers inherit the gap at close.

Evidence: Model-risk-management register with risk tier, validation schedule, monitoring metrics and retirement criteria

Wave 3 — confirmatory

AI/ML advisor

Deal team

Price impact

Infrastructure Layer (0 of 11 received)

GPU contracts, capacity guarantees, cloud concentration, committed-spend obligations and BIS export-control posture — the one layer where the buyer inherits a liability that does not appear in reported cost of goods sold.

- ☐ **Read every GPU contract for committed capacity, allocation guarantees, force-majeure, termination triggers, pricing and renewal mechanics.**

Committed capacity well above peak demand scores 5; the contract terms are what decide whether the growth plan is compute-supportable at all.

Evidence: GPU reservation contracts

Wave 3 — confirmatory

AI/ML advisor

Deal team

Price impact

- ☐ **Flag best-effort-only hyperscaler allocation with no committed-capacity SLA.**

It is the most common Infrastructure finding of 2025-2026 and becomes a continuity risk in capacity-constrained quarters; best-effort-only scores 1-2.

Evidence: Allocation terms in the cloud or GPU agreement

Wave 3 — confirmatory

AI/ML advisor

Price impact

- ☐ **Get a capacity commitment letter from each GPU provider covering the next 18 months.**

It is the buyer's evidence that compute supports the plan through close and the first 12-18 months post-close; no letter scores 1-2.

Evidence: Capacity commitment letters

Wave 3 — confirmatory

AI/ML advisor

Lender

Price impact

- ☐ **Compute the cloud concentration ratio — what share of compute spend runs on a single hyperscaler.**

Below 50 percent scores 5, above 70 percent scores 2-3, above 90 percent scores 1 — and the 70-percent line is our own working rule of thumb, not a published threshold.

Evidence: Cloud-spend breakdown by provider

Wave 2 — LOI

AI/ML advisor

Deal team

Price impact

- ☐ **Count how many model providers are in production and whether any fallback has actually been tested.**

ICONIQ found builders running around 3.1 providers on average precisely to hedge this; a single provider with no tested fallback concentrates a repriceable cost line.

Evidence: Provider routing policy and fallback test records

Wave 2 — LOI

AI/ML advisor

Price impact

[source](#)

- ☐ **Pull every cloud and GPU minimum-commit or take-or-pay schedule with remaining term, unused-commitment balance and true-up mechanics.**

These are off-P&L obligations in substance: the spend floor is owed whether or not usage arrives, and it rarely appears in reported cost of goods sold.

Evidence: Committed-spend schedules with unused-commitment balances

Wave 2 — LOI

Commercial counsel

Lender

Price impact

- ☐ **Ask what each committed-spend obligation does on change of control — accelerate, terminate with penalty, or stay silent and negotiable.**

It is the question asked last and mattering most; an acceleration clause lands as a cash item at close.

Evidence: Change-of-control clauses in compute agreements

Wave 3 — confirmatory

Commercial counsel

Lender

Price impact

- ☐ **Model the unused commitment balance as a deal liability rather than as a run-rate cost.**

A multi-year take-or-pay sized to a growth plan the target has already missed scores 1-2 and is a buyer-inherited liability, not an operating line.

Evidence: Utilization history against each committed-spend schedule

Wave 3 — confirmatory

Commercial counsel

Lender

Price impact

- ☐ **Request the disaster-recovery test log with scope, date, RTO achieved, RPO achieved and gaps identified.**

Quarterly tests with RTO under four hours score 5, annual tests score 3, no documented testing scores 1-2.

Evidence: DR test log

Wave 3 — confirmatory

AI/ML advisor

Confirmatory

- ☐ **Request the BIS export-control classification record for any AI technology or model exported, plus any licenses required.**

The January 2026 rule moving H200- and MI325X-equivalent chips to case-by-case review made export posture a deal-by-deal modeling exercise; no record scores 1-2.

Evidence: BIS classification record and foreign-affiliate screening program

Wave 3 — confirmatory

Regulatory counsel

Deal team

Price impact

- ☐ **Diligence any close on or after 10 November 2026 as if the BIS Affiliates Rule is live again.**

The stay runs one year from 10 November 2025 and the 50 percent ownership provisions return on 10 November 2026 unless BIS extends it.

Evidence: Affiliate-ownership screening analysis dated to the expected close

Wave 3 — confirmatory

Deal team

Price impact

[source](#)

Output Layer (0 of 9 received)

FTC AI-claims exposure, hallucination liability, customer indemnification posture and the insurance rider — the marketing-claim posture test is the most material because it determines the enforcement exposure surface.

- ☐ **Review every public marketing claim about AI capability, accuracy or substitutability against a substantiation file for that claim.**

Conservative claims with counsel sign-off and substantiation score 5; aggressive claims with no substantiation files score 1-2 under Operation AI Comply.

Evidence: Marketing-claims register with substantiation files

Wave 3 — confirmatory

Regulatory counsel

IP counsel

Price impact

- ☐ **Flag professional-substitution claims of the DoNotPay type — the February 2025 final order required \$193,000 plus subscriber notice.**

It is the FTC's clearest priced precedent for an AI capability claim and sets the substantiation bar the register has to clear.

Evidence: Claims register entries asserting professional equivalence, with evidence files

Wave 3 — confirmatory

Regulatory counsel

Price impact

- ☐ **Screen any facial-recognition or biometric-ID deployment against the FTC Rite Aid five-year ban running through approximately December 2028.**

The ban period sits inside typical deal cycles, and biometric-ID use is also an EU AI Act tier question, so the two findings compound.

Evidence: Biometric use-case inventory and any consent-decree obligations

Wave 2 — LOI

Regulatory counsel

Walk risk

- ☐ **Review the customer-complaint log for complaints alleging AI inaccuracy, bias, harm or substitution failure.**

Low volume with prompt resolution scores 5, a pattern-of-issue scores 2-3, and no log at all scores 1.

Evidence: Customer-complaint log with source, date, substance, response and resolution

Wave 3 — confirmatory

Regulatory counsel

Confirmatory

- ☐ **Request the hallucination-incident database with customer impact, root-cause analysis, mitigation deployed and recurrence indicator.**

It is the operational evidence that the target treats false output as a tracked liability; no formal tracking scores 1-2.

Evidence: Hallucination-incident database with review cadence

Wave 3 — confirmatory

AI/ML advisor

Regulatory counsel

Price impact

- ☐ **Read every material customer contract for AI-output indemnification, liability-cap carve-outs and the right to disclaim AI output.**

Open-ended AI-output indemnification is an unpriced tail the buyer inherits; a consistent posture across the base scores 5.

Evidence: Customer-contract indemnification index

Wave 3 — confirmatory

Regulatory counsel

Commercial counsel

IP counsel

Price impact

- ☐ **Check whether hallucination carve-outs are tied to documented model limitations rather than left open-ended.**

Mature targets carve out hallucinations attributable to documented limitations; inconsistent posture across contracts scores 2-3.

Evidence: Clause-level extraction of AI indemnity and carve-out language

Wave 3 — confirmatory

Regulatory counsel

Commercial counsel

Price impact

- ☐ **Request the CGL, errors-and-omissions and cyber policies with their AI riders, and have counsel read the rider scope.**

Insurers added AI-specific exclusions as well as riders and the terms vary widely; no rider on material policies scores 1-2.

Evidence: Insurance policies and AI rider endorsements

Wave 3 — confirmatory

Regulatory counsel

Commercial counsel

Price impact

- ☐ **Treat the December 2025 Rytr set-aside as a shift in FTC enforcement priorities, not a withdrawal of Operation AI Comply.**

The underlying risk surface is still the marketing-claim register and the substantiation files, so do not discount the layer on the set-aside.

Evidence: Regulatory-posture memo covering current FTC AI enforcement

Wave 3 — confirmatory

Regulatory counsel

Confirmatory

Governance Layer (0 of 16 received)

EU AI Act tier classification, the conformity assessment file, the NIST AI RMF program and red-teaming cadence — no AI risk register and no documented red-team cadence is the single most common Tier 1 finding.

- ☐ **Open at teaser stage with a one-page AI exposure summary: the five layer scores plus the EU AI Act tier classification summary.**

It is wave one of the disclosure model and lets a buyer price the AI risk before the seller exposes a single manifest.

Evidence: One-page AI exposure summary

Wave 1 — teaser

Deal team

Confirmatory

- ☐ **Sum the five layer scores into the 5-25 AI Deal Health Score and act on the band: 5-10 walk or asset-only, 11-17 price-chip, 18-25 healthy.**

The band is the decision, not the memo — below 11 the AI-specific indemnification exposure can exceed deal value.

Evidence: Completed 5-Layer scoring sheet with the composite score

Wave 1 — teaser

Deal team

Walk risk

- ☐ **Require an EU AI Act tier classification for every use case the target's AI deploys against — typically 3 to 12 per target.**

Tier classification belongs in the LOI wave because Governance is bimodal: days if the file exists, months if it has to be built.

Evidence: Per-use-case EU AI Act tier classification reviewed by EU counsel

Wave 2 — LOI

EU AI Act counsel

Deal team

Price impact

- ☐ **Walk any Unacceptable-tier use case for an EU-touching buyer — the Article 5 prohibitions have applied since 2 February 2025.**

There is no readiness reserve that prices a prohibited practice; it is a walk or a carve-out of that business line.

Evidence: Use-case inventory screened against the Article 5 prohibitions

Wave 2 — LOI

EU AI Act counsel

Deal team

Walk risk

[source](#)

- ☐ **For every High-Risk Annex III use case, request the conformity assessment file against the 2 December 2027 application date.**

The Digital Omnibus deferred Annex III from 2 August 2026 to 2 December 2027, but that date still sits inside a typical hold period.

Evidence: Conformity assessment file

Wave 3 — confirmatory

EU AI Act counsel

Deal team

Price impact [source](#)

- ☐ **Check the conformity file for Annex IV technical documentation, data governance, post-market monitoring, human oversight and cybersecurity.**

A complete file scores 5, a partial file 2-3, and none with a close near the deadline scores 1 and carries the full readiness chip.

Evidence: Annex IV technical documentation and the EU declaration of conformity

Wave 3 — confirmatory

EU AI Act counsel

Price impact

- ☐ **For Limited-Risk use cases, request the Article 50 transparency disclosure evidence — those obligations apply from 2 August 2026.**

Article 50 was not deferred; only the Article 50(2) marking duty, and only for systems placed on the market before 2 August 2026, has a transitional to 2 December 2026.

Evidence: Article 50 transparency disclosures per product surface

Wave 3 — confirmatory

EU AI Act counsel

Deal team

Confirmatory [source](#)

- ☐ **Size Article 99 exposure against total worldwide annual turnover of the post-close group, not the target's EU revenue.**

Fines of up to EUR 35 million or 7 percent for prohibited practices have applied since 2 August 2025 and are group-wide after close; Article 99(6) applies the lower of the two to SMEs.

Evidence: Article 99 exposure model against combined group turnover

Wave 2 — LOI

EU AI Act counsel

Deal team

Price impact [source](#)

- ☐ **For a target that builds or distributes a GPAI model, confirm the model card, training-data summary and copyright-compliance policy exist.**

GPAI obligations have been in force since 2 August 2025, so this is a live compliance gap rather than a future one.

Evidence: GPAI documentation set: model card, training-data summary, copyright policy

Wave 3 — confirmatory

Deal team

AI/ML advisor

Price impact

- ☐ **Map the risk-management program to NIST AI RMF 1.0 and the Generative AI Profile, with owners and cadence per Govern, Map, Measure, Manage.**

Documented alignment scores 5, selective alignment 3, none 1-2 — and the GenAI Profile adds confabulation, privacy, bias and misuse coverage.

Evidence: NIST AI RMF program documentation with function-level procedures

Wave 3 — confirmatory

Deal team

AI/ML advisor

Confirmatory

[source](#)

- ☐ **Read the ISO/IEC 42001:2023 certificate's scope statement and call the named certification body — do not score the certificate's existence.**

42001 is certifiable and the scope statement is where a narrow certificate stops covering the product — and no management-system certificate substitutes for an EU AI Act conformity assessment.

Evidence: ISO/IEC 42001 certificate with scope statement and certification body

Wave 3 — confirmatory

Deal team

Price impact

- ☐ **Request the red-team log with cadence, scope, findings and remediation — quarterly minimum for production GenAI.**

Quarterly with remediation tracking scores 5, annual scores 3, none scores 1-2, and the absence is half of the most common Tier 1 finding.

Evidence: Red-team log with test scope (prompt injection, jailbreak resistance, data exfiltration, bias, harmful content) and remediation tracking

Wave 3 — confirmatory

AI/ML advisor

Regulatory counsel

Price impact

- ☐ **Request the AI risk register with each risk's owner, mitigation, residual risk score and review cadence.**

No AI risk register plus no red-teaming cadence is the single most common Tier 1 Governance finding and triggers an automatic score of 1-2.

Evidence: AI risk register with executive review records

Wave 2 — LOI

Deal team

Data-protection counsel

Price impact

- ☐ **Request the document assigning AI diligence oversight to named senior management and the board, plus written risk-tolerance thresholds.**

A policy without thresholds is a brochure; the named-owner document is what turns a responsible-AI claim into an auditable process.

Evidence: Board and senior-management AI oversight assignment plus risk-tolerance thresholds

Wave 3 — confirmatory

Deal team

Confirmatory

[source](#)

- ☐ **Request the testing, evaluation, verification and validation record across releases — the trend in measured performance, not a snapshot.**

OECD steps two and four turn a vague claim into a document set, and a declining trend across releases is a finding a snapshot hides.

Evidence: TEVV record: test sets, metrics, tooling and measured performance across releases

Wave 3 — confirmatory

AI/ML advisor

Confirmatory

[source](#)

- ☐ **Request the incident log covering significant incidents, errors and safeguard-circumvention attempts, plus the stakeholder remedy channel.**

OECD steps five and six close the loop: without a remedy channel the target has no route to cure an incident short of litigation.

Evidence: Incident log and stakeholder remedy channel documentation

Wave 3 — confirmatory

Regulatory counsel

Confirmatory

[source](#)

Customer Contracts (0 of 8 received)

The outbound question most AI diligence skips: whether the target's own customer contracts actually permit training on the data the valuation model treats as a compounding moat.

- ☐ **Run clause extraction across the full customer contract set with a citation back to the source paragraph for every hit.**

The data asset you are paying for may be contractually unusable after close, and no credible published prevalence stat exists — only your own count with citations is defensible.

Evidence: Clause-extraction output with source-paragraph citations for every contract

Wave 3 — confirmatory

Deal team

IP counsel

Price impact

- ☐ **Count express training grants and check whether each survives assignment and names model training explicitly.**

It is the cleanest case, but a grant that does not survive assignment evaporates at close exactly when the buyer starts relying on it.

Evidence: MSAs, order forms and product terms carrying a training grant

Wave 3 — confirmatory

IP counsel

Deal team

Price impact

- ☐ **Count contracts silent on training rights and treat silence as the weakest position, not as permission.**

Silence is jurisdiction- and regulator-dependent and is the hardest position to defend to an acquirer's counsel.

Evidence: Older MSAs and self-serve click-through terms

Wave 3 — confirmatory

IP counsel

Price impact

- ☐ **Count no-training and no-derivative clauses, then ask whether any model was already trained on data covered by one.**

The data is unusable as a training asset, and a model already trained on it is a live exposure rather than a forward restriction.

Evidence: Enterprise-negotiated MSAs and security riders

Wave 3 — confirmatory

IP counsel

Deal team

Walk risk

- ☐ **Check DPA processor-only restrictions: training is not a documented instruction even where the MSA is silent.**

It closes the gap sellers use to argue silence equals permission, and it moves whole cohorts of data out of the asset.

Evidence: Data processing agreements with annexes and schedules

Wave 3 — confirmatory

Data-protection counsel

Price impact

- ☐ **Count deletion-on-termination clauses and ask for the unlearning or retraining plan behind each one.**

Deleting from a corpus does not delete from a trained model, so the clause creates an obligation the target may not be able to perform.

Evidence: MSA termination articles and DPA deletion terms

Wave 3 — confirmatory

Data-protection counsel

IP counsel

Price impact

- ☐ **Count change-of-control consent and termination rights separately from the training-rights count.**

One clause can strip the data right and the revenue at the same time, so the two exposures must be sized independently.

Evidence: MSA assignment articles

Wave 3 — confirmatory

Deal team

IP counsel

Walk risk

- ☐ **Test the chain of promises from customer to target to sub-processors to the upstream model provider's default tier.**

A target promising customers it never trains on their data while routing to a provider tier where that is not the default is the highest-value finding here.

Evidence: Customer no-training commitments matched against dated sub-processor and provider terms

Wave 3 — confirmatory

Data-protection counsel

IP counsel

Price impact

Research Team (0 of 8 received)

In an AI target the research team frequently is the asset, and the 2024-2025 license-and-hire run is the market pricing the research team rather than the code.

- ☐ **Ask who specifically trained each production model and whether they are still employed today.**

In an AI target the research team frequently is the asset, so a departed trainer is a value question, not an HR one.

Evidence: Per-model training-owner list with current employment status

Wave 3 — confirmatory

AI/ML advisor

Deal team

Walk risk

- ☐ **Ask how many current employees could retrain the flagship model from scratch — if the honest answer is one, that is the finding.**

A bus factor of one on the flagship model is the AI-specific version of key-person risk and drives the whole retention structure.

Evidence: Named list of engineers who could reproduce a full training run

Wave 3 — confirmatory

AI/ML advisor

Deal team

Walk risk

- ☐ **Ask whether the training pipeline is documented well enough for a new hire to reproduce a run, or whether it is tribal knowledge.**

Reproducibility is the difference between buying a capability and buying five people who may leave.

Evidence: Training-pipeline documentation and a reproduced run log

Wave 3 — confirmatory

AI/ML advisor

Price impact

- ☐ **Map every researcher's vesting cliff against the expected close date.**

Cliffs landing just after close convert the retention pool from an incentive into a payout with no lock-up.

Evidence: Cap table and equity vesting schedules for the research team

Wave 3 — confirmatory

Deal team

Price impact

- ☐ **Test whether non-solicits and non-competes are enforceable in the jurisdictions where the researchers actually live.**

An unenforceable restrictive covenant means the retention package is the only thing holding the asset in place.

Evidence: Employment agreements with jurisdiction-by-jurisdiction enforceability analysis

Wave 3 — confirmatory

Deal team

IP counsel

Price impact

- ☐ **Price what it would cost to replace this team on the open market, using the license-and-hire comparables rather than the code.**

Buyers repeatedly paid nine and ten figures for a few named researchers plus a non-exclusive licence — these were licence-and-hire deals, not acquisitions, and that is the replacement benchmark.

Evidence: Comparable license-and-hire transactions with named-researcher counts

Wave 3 — confirmatory

Deal team

Price impact

[source](#)

- ☐ **Structure retention pools, earn-outs tied to named individuals and staged consideration off that replacement number.**

It converts a key-person finding into deal mechanics instead of a memo paragraph nobody prices.

Evidence: Draft retention and earn-out schedules naming the individuals

Wave 3 — confirmatory

Deal team

Price impact

- ☐ **Run the IP chain upstream of the acquirer: license-backs and IP-assignment rights already held by your own investors and model partners.**

OpenAI's reported roughly \$3 billion Windsurf deal collapsed on a term in the buyer's own arrangements, not on anything in the target.

Evidence: Buyer-side IP-assignment and license-back terms with investors and model partners

Wave 2 — LOI

IP counsel

Deal team

Walk risk

[source](#)

Inference Economics (0 of 13 received)

Recompute gross margin with inference inside cost of goods sold, by cohort, and then ask who controls the inputs to that line.

- ☐ **Rebuild gross margin with inference inside cost of goods sold before applying any multiple.**

The most common restatement in AI diligence is model spend booked in R&D, which flatters gross margin by anything from a few points to a category; resistance to reclassifying it is the tell.

Evidence: Trial-balance mapping for model spend and the R&D-to-COGS reclassification schedule

Wave 2 — LOI

Deal team

Lender

Price impact

- ☐ **Compute inference cost per active account by cohort and by plan from provider invoices, and flag cost rising faster than price.**

Cost per account outrunning price per account is a margin trajectory problem no revenue growth rate fixes.

Evidence: Provider invoices divided by monthly active accounts, by cohort and plan

Wave 3 — confirmatory

Deal team

AI/ML advisor

Price impact

- ☐ **Pull API logs tagged by model and product surface, and flag frontier-tier models serving free-tier or trial traffic with no routing policy.**

It is the cheapest margin fix a buyer can underwrite, and its absence signals the cost line is not actively managed.

Evidence: API call logs broken out by model and by product surface

Wave 3 — confirmatory

AI/ML advisor

Price impact

- ☐ **Compute third-party model spend as a share of cost of goods sold and flag any single provider above half of COGS with no tested fallback.**

It measures how much of the margin a supplier can reprice unilaterally, which is the whole wrapper question in one number.

Evidence: Vendor concentration schedule inside cost of goods sold

Wave 2 — LOI

Deal team

Lender

Price impact

- ☐ **Ask what share of production inference calls terminate at a third-party hosted model.**

It is the first of the three wrapper numbers and it sets the pricing category — margin-and-switching-risk buy, or technology buy.

Evidence: Inference routing breakdown by destination model

Wave 2 — LOI

AI/ML advisor

Deal team

Price impact

- ☐ **Ask what the target still owns after a model swap: distribution, workflow depth, proprietary corpus, evaluation infrastructure.**

A wrapper is not a bad buy, it is a different buy — what survives a model swap is what you are actually paying for.

Evidence: Asset inventory of what persists through a base-model change

Wave 2 — LOI

AI/ML advisor

Deal team

Price impact

- ☐ **Request the model-deprecation playbook: what engineering actually did the last time a provider retired an endpoint or changed a rate limit.**

It is the highest-information document in the wrapper test because it shows behaviour rather than intent.

Evidence: Model-deprecation playbook plus the incident record of the last migration

Wave 3 — confirmatory

AI/ML advisor

Price impact

- ☐ **Request provider invoices for the last eight quarters together with any contract carrying rate commitments.**

Invoices are the evidence base for every other number in this layer; the pitch deck is not.

Evidence: Eight quarters of provider invoices and rate-commitment contracts

Wave 3 — confirmatory

Deal team

AI/ML advisor

Confirmatory

- ☐ **Test the margin plan against realized cost per unit of work, not against list-price deflation.**

Inference prices fall fast but unevenly — Epoch measured 9x to 900x per year, frontier workloads sit at the slow end, and Epoch cautions the fastest drops may not persist.

Evidence: Realized cost-per-unit-of-work trend by workload

Wave 3 — confirmatory

Deal team

AI/ML advisor

Price impact

[source](#)

- ☐ **Read the target's gross margin against Bessemer's archetype spread — about 25 percent at one archetype and roughly 60 percent at another.**

Two AI companies at identical revenue can sit 35 gross-margin points apart, so the spread is the ranking tool, not a valuation input.

Evidence: AI-inclusive gross margin computed on the rebuilt bridge

Wave 3 — confirmatory

Deal team

Confirmatory

[source](#)

- ☐ **Refuse any published dollar range or week-count for AI diligence — no independent primary benchmark exists for either.**

Every specific figure on this topic traces to vendor content marketing with no sample, methodology or collection date, and a budget memo cannot defend it.

Evidence: Scope memo stating the absence of a published benchmark

Buyer-side scoping

Deal team

Confirmatory

- ☐ **Bound the AI workstream's cost from corpus size and messiness, contract volume, regulatory surface, specialist count and seller readiness.**

Those five drivers are what actually move the number, and they price the AI workstream as a premium on top of technology and IT diligence.

Evidence: Scoping estimate built from the five cost drivers on this deal

Buyer-side scoping

Deal team

Confirmatory

- ☐ **Plan the timeline by phase: Data Layer two to four weeks as the long pole, Model, Infrastructure and Output concurrent, Governance bimodal.**

Governance is days if a conformity assessment file and risk register exist and months if they must be built, which is why tier classification belongs at LOI.

Evidence: Phase-level workplan with the document-production gates named

Buyer-side scoping

Deal team

Confirmatory

The three disclosure waves

- Buyer-side scoping — planning steps the deal team takes before any request goes to the seller; not a seller disclosure wave.
- Wave 1 — teaser stage: a one-page AI exposure summary (the five layer scores plus the EU AI Act tier classification summary) before any deep folder is visible.
- Wave 2 — LOI stage: the AI risk register, the model inventory, the compute-spend picture and the EU AI Act tier classification in full.

- Wave 3 — confirmatory diligence: every layer folder in full, including the training-data manifest, the conformity assessment file and the red-team log.

Every check is a projection of [our AI due diligence guide](#) — no figure appears here that is not in it, and the score-band and concentration rules of thumb it labels as ours keep that label. Ticks are stored in your own browser only; nothing you check is sent to or stored on our servers.

Source: [peony.ink/blog/ai-due-diligence](https://www.peony.ink/blog/ai-due-diligence) · <https://www.peony.ink/tools/ai-due-diligence-checklist>

Print view: all 90 checks expanded.