

How to diligence B2B SaaS for AI

Since the emergence of ChatGPT in late 2022 and the concomitant rapid evolution in AI capabilities, AI has become the dominant lens through which private equity investors assess software businesses. Read more about how Riplo helps investors assess AI risks and opportunities for software deals.

In this paper

Introduction

01 Market structure

02 Competitive dynamics

03 Value creation

04 AI capabilities

Riplo's view


Introduction

AI has become the pre-dominant theme in private market software investing, driving deal teams to add dedicated AI impact diligence (AI impact DD) workstreams or embedded AI assessments into CDD and TDD as part of the diligence process for software assets. For these workstreams, their first port of call has often been their existing commercial or technical advisors and early efforts to formalise AI impact DD have largely been shaped by relationships with existing advisors.

Most traditional advisors have turned to the same backward-looking methodologies they rely on for other diligence topics: a combination of desk-based research, survey work and expert/customer calls. This works well when market participants have a good grasp of what is changing and how quickly, when customer expectations, the competitive set and required product features are well understood by potential targets as well as their customers, and, perhaps most importantly, when market and technology conditions are stable enough for an asset's recent performance to serve as a useful guide for an asset's trajectory over the average hold period, i.e. around five to seven years.

However, these backward-looking methodologies fail to capture AI risks and opportunities over a 5–7 year hold period, given the recent pace of progress in frontier AI and the heterogeneity in adoption across industries. The shift from ChatGPT's launch in November 2022 to Claude Cowork's multi-agent capabilities in February 2026 is a good illustration of how quickly the technical frontier can move.

Evaluating a 2028 exit based on 2023 technology and market environment has proven obsolete within months, and the same remains true today: rapid improvements in model capabilities continue to re-shape the commercial and technical environment in which (software) businesses operate.



Similarly, market participants and traditional vendors are generally not well versed in AI's trajectory and frontier capabilities. Consequently, as AI has become the main driver for change in B2B SaaS, they struggle to project AI capabilities into medium-term commercial implications for their industries. This is true across diligence topics: from disruption risk and pricing, to retention/renewal, feature importance rankings and technical discussions on moat or barriers to entry. Put simply, asking a CFO or procurement manager about budget evolution or likelihood to switch for e.g. accounting software over the next 12 months, the quality of their answer will be limited by their grasp of emerging frontier AI capabilities, and concomitant commercial offerings that may become available over the next 3–5 years, not expertise in their specific field.

Closing this gap requires a new approach: a frontier-oriented, forward-looking AI diligence focused on AI's commercial implications for the target's market, the target's competitive position and its AI capabilities. This means evaluating whether and when high-value use cases will become available in a given industry (technology lens), what their commercial impact will be (commercial lens), how long it will take for these use cases to get adopted, what this means for the asset's competitive position and right to win, and what it takes for a given asset to execute towards this, given their existing capabilities. Differently put, AI impact diligence revisits foundational investment questions across four core dimensions:

- **Market structure:** How defensible is the asset's market and business model as AI improves?
- **Competitive dynamics:** How does AI affect competitive dynamics and shift the right to win?
- **Value creation:** What are opportunities from AI, both for revenue and efficiency gains?
- **AI capabilities:** Does the asset have a credible vision on AI, and the resources and track record to execute against it?

Drawing on lessons from past deals, this paper presents a practical approach on how to do so.

01

Market structure

How defensible is the asset's market and business model as AI improves?

What ICs are asking:

How much of an impact does AI have on the asset's industry? Does the market remain defensible as AI improves? What shape might new commercial offerings take and how long do incumbents have to respond? Which structural features of the market protect incumbent revenues and how durable are these properties?

How defensible is the asset's market and business model as AI improves?

1.1 Understand level of AI impact on target industry

To form a view of the depth of AI's impact on a given target (market), investors need to quantify AI's potential to complete or complement human workflows across core user personas. This requires decomposing adjacent end-to-end workflows into discrete steps, identifying the specific user persona for each, and then evaluating time spent, task complexity, and automation potential. Generally, workflows that are repetitive/high volume, pattern-based with tolerance (i.e. low cost) for mistakes and have clear success/fail criteria are better suited for AI automation, whereas deterministic, audited and/or judgement-based workflows are less so.

Table 1: Indicators of AI potential by workflow

Lower AI potential	Higher AI potential
Deterministic, requiring reproducible outputs	Repetitive, high-volume work with limited human intervention
Regulated, with liability if outputs are inaccurate	Pattern and transformation-based work: mapping, classification, extraction
Audited, requiring a traceable trail for each decision	Tolerance for probabilistic output, low cost for errors
Multiparty, requiring handoffs between organisations	First-pass work subject to human review before use

While AI-native challengers might benchmark AI-induced time savings against a fully manual process, a more robust analysis compares end-to-end time and cost of real customers' existing workflows (including any incumbent software) with a "full potential" AI-native workflow, considering integration, workflow redesign, retraining, and other adoption costs as well as operating costs.

Cross-referencing the quantitative perspective against KPCs (Key Purchasing Criteria) reveals why, whether and how deeply buyers and potential customers care. AI-native solutions typically impact purchasing decisions where speed, cost, and ease of use are top of mind, but hold less immediate weight when buyers prioritise domain knowledge, vendor reputation, end-to-end coverage, or implementation support. That said, KPCs themselves evolve as AI capabilities mature, and different product bundles develop and features move from differentiators to table stakes. Governance, auditability and validation features, for instance, increase in relevance, as AI assumes an increasing share of work.

1.2 Identify the shapes of new commercial offerings

Mapping this view against the target's products and revenue lines provides both a bottom-up view of AI exposure, and an indicative picture on where AI-native product offerings might emerge (either via competitors or substitutes). Where AI impact is high, stand-alone AI-native products likely emerge to capture newly expanded Total Addressable Market (TAM). Where AI impact is medium, i.e. AI provides real but incremental value, AI features are likely to be monetised via feature bundling, tiered pricing, or as add-on modules (via cross-sell/upsell). Where AI impact is low, AI rarely shifts commercial dynamics, i.e. AI capabilities quickly commoditise into table-stakes functionality with minimal pricing power.

Table 2: Shapes of new AI commercial offerings by AI disruption risk

AI impact	Definition	Commercial expression
High	AI unlocks capabilities existing vendors could not previously provide, changing what customers purchase; customers care and are willing to switch.	TAM expansion: new products and commercial offerings
Medium	AI adds real but incremental value; not enough to make customers switch; human control points remain the same.	Cross-sell / upsell: new pricing and additional modules
Low	AI becomes a commodity capability customers expect as standard.	Commodity features: table stakes (no pricing impact; maybe cost pass-through)

Riplo's perspective:

High structural barriers delay, rather than prevent AI disruption. Incumbents (or their commercial advisors) often treat complex workflows, switching costs and other structural barriers as permanent moats. In reality, if AI is sufficiently transformative for the underlying workflow, customers will eventually switch if incumbents don't adapt. Diligence must clearly articulate the difference between sustainable defensibility and more transitory barriers that delay disruption over the holding period.

Example:

Workday's moat is its position as system of record. While frontier models can easily complete many of the individual tasks in Workday's remit (e.g. generate performance reviews, process expenses, or schedule shifts), Workday retains structural defensibility because it also serves as the source of truth for payroll routing, regulatory compliance logic, and complex, multi-entity data structures. As AI-native entrants (or native foundation models) might capture generative workflows, enterprise buyers will not swap out core ledgers of employee records for superior AI interfaces alone.

In this context, potential (un-)bundling across end-to-end workflows is a critical and often overlooked factor. Software revenue is most vulnerable when customers can isolate automatable tasks from broader core workflows. By implication, integrated platforms and systems of record remain defensible if automatable steps cannot be cleanly separated. However, if buyer preference tilts toward consolidation across workflows (i.e. reduce the number of overall tools required in or across a certain (number of) domain(s) in their business, e.g. in HR/payroll software), AI-assisted software development might encourage larger, adjacent platform players to enter niche adjacent spaces and "wrap" specialized point solutions into broader enterprise suites (e.g. expenses within payroll or broader HR platforms).

1.3 Speed of change. Understand barriers to new development & adoption of new solutions

While the analyses above quantify the scale of AI's impact, structural barriers to entry determine the pace at which these shifts appear, and who is best placed to benefit from them. Such barriers generally benefit incumbents, by increasing complexity, cost or time to develop new AI products and/or slow down adoption of net-new (AI) products (See Table 3 below).

In summary, the impact of AI on the associated workflows determines how profoundly an industry is going to shift, whereas the barriers help understand how quickly such changes are going to occur. Together, they define how much (and urgently) incumbents need to adapt.

Table 3: Barriers to AI adoption

Dimension	Type of barrier	Why it slows adoption	Evidence to consider
Regulatory requirements	Harder to develop net new product	Compliance requirements restrict entry	Filings, audit trails, certification requirements
Domain knowledge	Harder to develop net new product	Product encodes industry rules and edge cases	Edge-case rules, format libraries, evaluation data
Proprietary data	Harder to develop net-new product	Exclusive access to customer or industry data	Data ownership terms, scale of data assets
Workflow integration	Harder to develop net new product + harder to switch	Product is embedded in customers' existing systems	Integration depth, downstream dependencies, number of systems integrated in
Technology complexity	Harder to develop new product	Specialist talent, time; execution risk	Software and AI complexity, scale, data volumes and latency requirements
Switching costs	Harder to switch	Replacing the product is costly, disruptive, or risky for business continuity	Cohort retention, products per customer, system/data and workflow integration
Distribution	Harder to switch	Owned customer channels, supported by brand trust, take time to re-build	Win/loss rate against competitors, RFP presence
Budget & contract cycles	Harder to switch	Limited share of contracts coming to market in a given time, slowing adoption	Contract terms and cycle lengths

Competitive dynamics

How does AI affect competitive dynamics?

What ICs are asking:

How does AI affect the asset's market share and right to win? Are new competitive threats emerging? Who is best placed to capture value from the shifts caused by AI? Where do threats come from and how is the asset positioned to defend against them? Could AI disrupt the category of the target's customers?

How does AI affect competitive dynamics?

The market assessment above builds the foundation for (1.1) where (new) value accrues, what shape new commercial offerings may take and which parts of the workflow they might focus on; as well as (1.2) the expected pace of adoption & adaptation in the market.

2.1 Map the full landscape of potential competitive threats

Competitive threats could emerge from several different angles. Analysing how each could enter, including a review of their capabilities and incentives to do so helps sharpen the view on which competitive threats are acute, and how well the asset is or should be positioned to defend its turf.

Table 4: Competitive threat landscape

Competitor type	Questions to answer to assess viability
Foundation-model or big-tech entry	Is the market large and strategically relevant enough to attract frontier-model or platform attention? Do foundation model labs have a clear reason to enter?
AI-native challengers	Have startups entered the space, and/or show evidence of customer adoption, live deployments, and RFP participation? Are ACVs, deal cycles and industry size large enough to be attractive?
Existing competitors	Does AI shift the asset's relative position against the current competitive set? How mature are competitors' AI capabilities and what is their strategic orientation (incl. ownership structure)?
Adjacent system provider	Could an upstream, downstream or adjacent provider enter the space and bundle the target's software within their own offering? Do potential providers have the relationships, data, and product foundations needed to expand? Do they tend to partner, acquire, or build if they chose to enter?
Customer in-housing	What would it take for customers to build in-house? Do they have the capabilities and strategic incentive to in-house, given total cost of ownership (TCO) across building, operating and maintaining such software themselves, and can leadership commit the sustained focus this requires without pulling it from other priorities?

Riplo's perspective:

"Vibe coded" replicas rarely pose the primary threat to B2B SaaS. AI-native challengers often lack distribution and trust, while in-house builds come with costs for hiring talent, managerial focus and maintenance, which particularly non-software businesses can rarely justify. Instead, the greatest risks come from incumbent competitors and adjacent platforms as they already own customer relationships, workflow and domain knowledge, with access to (superior) resources to execute over a multi-year horizon.

Example:

Thomson Reuters' distribution, proprietary data and trust shield core revenues from standalone AI models.

When Claude launched its legal capability plug-in in early 2026, public markets panicked, driving a 30%+ drop in Thomson Reuters' share price over the following months. Yet, Q1 revenues grew 10%. Thomson Reuters' moat is not their search UI, but institutional trust and proprietary data: legal and tax buyers require verified, source-traceable data, as errors carry risk of liability. The core diligence question for TR would be who (if anyone) is best placed to replace their role as trusted data vendor.

Mapping the competitive landscape helps understand who has the right to win. This requires looking beyond the traditional set of competitors and assessing a broader set of actual and potential threats by understanding both their capabilities (distribution, domain knowledge, data, capital, talent etc.) and incentives (strategic orientation and priorities, vision, ownership structure, incentives and execution speed).

2.2 Identify right to win and dive deeper on highest impact threats

Once high-priority competitive threats are identified, AI impact diligence must dive deeper into their strategic posture, incentives and capabilities e.g. by examining their org structure, AI talent density, AI posturing and ad spend, or targeted expert interviews with recently departed employees.

Counterintuitively, the single biggest impact of AI in B2B SaaS has been on software development itself, disproportionately benefiting foundation model labs and incumbents, not start-ups. Agentic coding has made it significantly faster and cheaper to build software, lowering the cost of entry for AI-native challengers. At the same time, incumbents (or companies planning to develop software in house) now have access to the very same tools, meaning development speed alone will not provide a durable edge for any one player. If anything, incumbents are better placed to use it, since they already hold the domain expertise and just need to translate it into code faster. The barrier is no longer writing software; it's capturing domain expertise into the product quickly, comprehensively and cheaply. Competitive advantage shifts increasingly to non-code assets, such as domain expertise, distribution networks, proprietary data or the ability to attract and retain talent. Conversely, incumbents in niche markets face increased risk from adjacent and potentially vastly out-scaled competitors motivated to enter their niche by reduced build costs.

Value creation

What are AI-enabled opportunities for revenue growth and efficiency uplift?

What are AI-enabled opportunities for revenue growth and efficiency uplift?

3.1 Identify & prioritise opportunities

AI creates varying opportunity for value creation across three core areas for software businesses: revenue from net-new products, opportunities for enhanced monetisation of existing products (including new features, cross/up-sell or pricing), and efficiencies in core processes in back-office and adjacent functions. AI diligence should identify and evaluate use cases across all three, balancing financial impact and feasibility against existing capabilities as well as time to value. Any AI initiatives should be clearly linked to their respective value levers.

Table 5: AI value creation opportunities by business function & time horizon

Area	Description	Diligence & operational approach
Short-term Operational efficiency (margin expansion)	Productivity gains across core product delivery (e.g. R&D) or G&A functions (e.g. sales, marketing, finance, customer support)	Build internally only where solution is a competitive differentiator and the economics of build and maintain justify it; buy OTS tooling, for non-differentiated back-office tasks to minimise maintenance & managerial distraction
Medium-term Existing product enhancement (monetisation uplift)	Revenue uplift driven by embedding AI capabilities into core modules (e.g. as add-ons, bundling, cross/up-sell, different packaging or value-based updates to pricing)	Leverage existing engineering capabilities and sales channels. Lower execution risk than net-new product development while driving more revenue growth.
Long-term New product revenue (TAM expansion)	Revenue from stand-alone, AI-native products that address the serviceable market	Evaluate 0-to-1 feasibility (commercially and technically). Highest risk; requires dedicated R&D resources, G2M approach and resource, and cross-functional collaboration across the org.

3.1.1 Net new product development

Net new product development is most promising for industries most impacted by AI but remains under-explored by private equity sponsors. As product development is inherently iterative, time consuming and complex and many PE-backed businesses lack the required capabilities, sponsors routinely default to short-term cost reduction or incremental feature builds over underwriting transformative revenue expansion.

Building agentic products can require 24–48 months of “zero to one” R&D, dedicated product discovery, iterative delivery, an overhaul of go-to-market from dedicated collateral and positioning, as well as training and incentivising the sales team, driving adoption and customer success, defining and training customer support, and all the change management efforts that come with it. If sponsors or management want to go down that path, diligence needs to assess rigorously whether the asset’s current management team, product and engineering organisations have the capabilities required and which additional capabilities are required.

3.1.2 Existing product monetisation

Instead of creating new stand-alone software, embedding AI capabilities directly into established products can increase utility for clients and capture value.

Pricing, packaging & tiering strategy: AI enhancements should not be given away as free platform updates. Diligence should evaluate packaging mechanisms, e.g. dedicated “Pro” or “Enterprise AI” tiers for modular add-ons, usage-based consumption models alongside price increases or up-sell for traditional seat-based licenses. As embedded AI features can create value for end users by automating (sub-)tasks, annual renewals create triggers for upsell while deepening workflow integration.

Low(er) execution risk: Because existing products rely on established go-to-market channels and customer relationships, execution risk remains significantly lower than launching net-new products, and can help an organisation validate internal capabilities and market sentiment more incrementally. On the flip side, sloppy execution might alienate existing customers, and too much focus on incremental innovation and monetisation widens the gap to more innovative competitors that invest in re-inventing customer workflows in a more AI-native way.

3.1.3 Operational efficiency & margin expansion

Many AI-enabled efficiency initiatives require minimal structural changes to core product architecture or organisational structure and can be executed within the first 6–12 months of ownership. Common levers for software include:

- **R&D productivity:** Embedding AI assistance into software and product workflows (e.g. automated code generation, testing, refactoring) reduces cycle times and increases throughput. Diligence should assess whether productivity uplift should be budgeted as cost savings or reinvested to accelerate feature delivery.
- **Professional services & time to value:** AI-driven ingestion, schema mapping, data migration and configuration compresses implementation cycles and shortens time-to-value, reduces professional services spend, and accelerates revenue recognition.
- **G&A and back-office automation:** High-volume, repeatable tasks across customer support, finance, legal, and sales make for clear automation targets. Automating level-1 support and streamlining e.g. contract analysis often expand margins without degrading service or NPS.

Efficiency initiatives carry the least risk, as metrics are (more) clear and measurable, technological feasibility often already validated in the market, and off-the-shelf solutions with dedicated implementation partners might already exist. While less uncertain than product innovation, full potential estimates for efficiency levers must be carefully adjusted for pilot performance and value realisation, with realistic timelines for adoption. Change management and senior sponsorship/management support are crucial to ensure value realisation.

Across all initiative types, quantification exercises should rely on the company's own operating and pilot data (e.g. on previous product or feature launches) wherever possible. External case studies and analyst estimates provide a cross-check but confidence in delivery should be adjusted accordingly.

Riplo's perspective:

AI revenue potential should move from narrative to the financial model. Rather than writing off any AI-related growth as “unproven”, deal teams should rigorously test each initiative against impact, cost and feasibility. Explicitly distinguishing efficiency, monetisation and growth opportunities with adjusted probabilities and returns, gives investment committees the confidence to price AI opportunity into the investment case.

Example:

Asset structure determines value creation levers. A Contract Lifecycle Management platform managing past data can leverage AI for aggressive TAM expansion by automating legal counsel workflows, whereas a lightweight contract generator/reviewer point solution is limited to margin protection. Access to the data layer unlocks value-based pricing and budget capture from adjacent human workflows; lacking it restricts AI gains to incremental cost reduction/efficiency plays.

3.2 Define a path ahead

Once the most promising levers are identified, AI impact diligence must bring clarity on the implementation approach, and contextualising AI initiatives within the broader value creation plan (VCP) to maximise value.

Build vs buy. Despite recent pessimism about B2B SaaS as a category, complex functional use cases are best bought off the shelf. Similarly, generic AI productivity tools have reached sufficient performance for most generic productivity workflows (e.g. drafting sales decks, updating CRM, basic data querying and reporting) and beating their performance may not be worth the (often sizable) effort of custom development or implementing dedicated software. Custom developments should only be considered if no off-the-shelf solutions are available, customisation warrants the extra effort, a broader strategic rationale justifies the higher TCO (total cost of ownership) and/or there is a high cost to waiting for the market to provide solutions. Generally, for undifferentiated back-office processes buying off-the-shelf software is superior, while R&D capabilities should focus on core product enhancements.

Total cost of ownership & maintenance. In the calculus above, an important (cost) factor to consider is the cost of maintenance of (agentic) systems after the initial build. Even seemingly “simple” automations (Claude skills etc.) require owners, governance, monitoring and regular maintenance to ensure they remain in sync with evolving business processes, often binding more expensive resources for longer than originally envisaged or planned for, with a concomitant negative impact on the associated business case.

Execution risk. For either type of initiative, value is only realised when employees or customers adopt the new product (features) and associated ways of working, which can often require both time and careful change management. All estimates should therefore carefully consider the asset's record of delivering on both software development and/or change management initiatives.

Contextualise within wider value creation plan. Finally, AI initiatives should be prioritised against the broader value creation plan as non-AI-based opportunities, such as M&A, commercial re-positioning or other strategic initiatives, which compete for the same budgets and management attention. If other options offer higher returns and easier implementation, AI initiatives should be prioritised and sequenced accordingly.

AI capabilities

Does the asset have a credible AI roadmap, and the resources to execute against it?

What ICs are asking:

How credible are the asset's AI plans and are they resourced appropriately? How do these compare against peers? What additional talent, technology, or investment might be required?

Does the asset have a credible AI roadmap, and the resources to execute against it?

Forming a view on AI capabilities of an organisation requires understanding "traditional" topics such as leadership, vision and roadmap, talent pools, existing tech and AI infrastructure and software development related process maturities. On top of this, capturing value from AI requires ability to successfully execute change, making a successful track record of change within the organisation indicative of the ability to benefit from AI.

4.1 Assess leadership, vision & roadmap

In the current exit environment, most leadership teams express AI-related ambitions. The depth of commitment and success to date can be pressure-tested with three considerations:

- **How well does leadership understand AI**, considering their backgrounds, public statements and previous experience, etc.?
- **How ambitious, feasible, and commercially viable** is the leadership's AI vision, and does it connect to specific, resourced initiatives that are tied to value? Are AI initiatives tied to EBITDA impact? (How) is this tracked? Are AI initiatives supported by appropriate budgets, resources and timelines?
- **What progress has already been made**, and did AI initiatives, resourcing and narrative pre-date the exit process? Which parts of the commercial and technological assumptions have already been validated?

4.2 Understand resourcing and talent

Deal teams should assess whether the asset has appropriate resources to deliver against the roadmap, by considering:

- **What is the depth and quantity of AI talent** and how appropriate are they for the requirements of the chosen use cases? Where is AI talent located? Do the appropriate roles, titles and career paths exist? Individual "Hero hires" might look good in the short term but retention can be challenging. If the plan requires additional resourcing, are all required roles budgeted, costed, and tied to realistic start dates? Insufficient capacity and timelines will lead to delays.
- **How is AI strategy & execution integrated into the organisation** and is this appropriate for the chosen use cases? Embedded, cross-functional teams, i.e. those building the product understand the industry well enough to encode its rules, workflows, and edge cases into the product, are superior to an isolated AI or innovation team, while insufficient domain expertise limits product quality.

Answers to these questions can help evaluate delivery confidence for revenue initiatives and build-versus-buy choices for efficiency initiatives.

4.3 Evaluate technology, data & architecture

Ability to execute against the roadmap requires assessing availability of data and associated rights, traditional product/tech infrastructure and existing AI infra capabilities. Investigating these topics clarifies asset's build vs buy decisions, integration requirements, investment needs and timelines to execute, and can give an indication of overall organisational AI maturity.

Table 6: Technology, data & architecture assessment

Dimension	Questions to answer & relevant evidence
Product & traditional tech architecture	What is the status of cloud migration (if applicable)? What is the API accessibility of data? What data/events can be streamed? How modular is the existing software? (monolith vs microservices) Evidence: Architecture diagrams and documentation
Data rights & readiness	What data is available? Do contracts permit use of customer and industry data, including training, fine-tuning, and evaluation where relevant? Evidence: Customer contract terms, data-ownership clauses, data samples, architecture & API documentation, data-governance policy & retention rules
AI eng infrastructure	What is the state and stack of models, data pipelines, evaluation tools, monitoring and test environments? Can these support reliable development? Is the model harness architecture appropriate for the respective use case(s)? Evidence: AI-specific tooling stack, harnesses and eval setup and measurements

4.4 Assess R&D processes and SDLC

Finally, deal teams should assess the asset's AI R&D processes across five areas, to understand whether they have been re-aligned to AI-native ways of working:

- **Product management:** How mature are processes to set product priorities, make trade-offs, and assign ownership? How are users, customers and wider business involved?
- **Product delivery:** How consistently does the team deliver against its roadmap? What does its release cadence, delays, and rollbacks show about past execution?
- **Agentic coding:** How (much) are AI tools used across the product lifecycle? What measurable effect have they had on delivery speed, test coverage, and code quality?
- **AI governance:** How well is the use of AI tools governed, with e.g. a tools register, a usage policy, and rules for how data is kept and deleted?
- **Documentation:** How well are the codebase, architecture, and technical decisions documented? How dependent is the organisation on knowledge held by single engineers?

Riplo's view

AI impact diligence requires a dedicated methodology.

AI impact diligence requires a dedicated methodology. Deal teams, ICs, and operating partners need confidence that an AI impact diligence partner forms an integrated perspective that bridges commercial judgement and technical expertise. Although some funds have started mandating dedicated AI impact diligences, many funds still rely on in-house teams or traditional vendors to cover AI topics as part of CDD or TDD, putting them at risk of missing what is increasingly the most important swing factor on any deal.

Table 7: Alternative approaches to AI impact diligence

AI impact diligence alternative	What they focus on	Where they fall short
Deal team	Deep understanding of the business, market sentiment and expert calls	Lack technical and frontier AI understanding, time constrained
Operating Partner	Expert perspective on client's AI stack and technology credibility	Lack of deep familiarity with the asset and cross-industry view of the art of the possible, time constrained
Chapter in CDD	Survey work or consulting firm expert judgement on market sentiment and disruption risk based on past data or next 12–18 months	Lack of technical and frontier AI understanding and depth, methodological bias towards the past, often treated as afterthought
Chapter in TDD	Technical view on "effort to replicate" core functionality, assessments of technical AI maturity of asset, disconnected from value	Lacks industry and business context to identify use cases or understand commercial implications; can be disconnected from frontier and not sufficiently forward looking

We would be delighted to work with you on your next deal.

About Riplo

Riplo provides AI-impact due diligence for private equity funds across Europe and the US, assessing assets' market position, defensibility, competitive position, value-creation opportunities, and capabilities through an AI lens. We are backed by Cherry Ventures and Blue Lion, alongside angels from McKinsey, BCG, Hg, OpenAI, Goldman Sachs and Apple.

Get in touch to discuss how Riplo can help on your next deal.

Tobias Haefele
CEO and Co-Founder
tobias@riplo.ai

Bhavna Dewan
Founding Director of Advisory
bhavna@riplo.ai

Benjamin Preisler
Chief of Staff
benjamin@riplo.ai