



The AI Studio Perspective

AI Changes What You Are Underwriting

Why AI value creation in private equity requires redesigning the asset, not just deploying better tools.

Private equity has a well-developed playbook for creating value: operational efficiency, pricing, sales effectiveness, procurement, technology modernization, add-on acquisitions, organizational redesign. These levers work because they rest on fairly stable assumptions about what it costs to build and run a business, and how much of the work requires people. AI unsettles those assumptions, which is why it sits awkwardly as one more item on the list. **A more useful starting point than asking where AI can be applied is to ask whether, designing the company today with what is now technically possible, we would design the same company. For many portfolio companies, the honest answer is no.**

Few companies are seen as software companies, but nearly all of them now run through software. Pricing, planning, customer service, reporting, operations and much of the product itself are encoded in systems, often bought from SaaS vendors rather than built in-house. Also, all AI tools and products are a software product, just not owned by you. AI changes what those systems can do, and it sharply lowers the cost of building new software, automation and decision-making capability. **That is why we treat AI transformation as an engineering problem, in the broad sense of the word: understanding how the company actually works, which constraints shape it, and how it should be rebuilt once those constraints shift.**

From tools to redesign

Most AI programs are still organized around tools and pilots, such as copilots, assistants and automated reports. These can, sometimes, pay for themselves, but they leave the company's production system largely as it was. It helps to separate three levels of intervention. Augmentation gives individuals better tools. Automation redesigns specific workflows. Re-engineering rebuilds the company around capabilities and economics that were not available before. Most organizations are still concentrated at the first two levels, while the largest opportunities increasingly sit at the third, where the questions are harder to fit into a traditional pilot plan:

- Which work should stay with people, which should become software, and which should be handled by agents?
- Which processes should be removed rather than automated?
- Which capabilities were never worth building at the old cost, but now are?
- Which services can become products, and which products can become systems that improve with use?

Working through these well takes more ambition than an adoption program, and a willingness to revisit long-held assumptions about headcount, margins and scale.

From activity to value

A company can have high AI adoption and still have captured almost no economic value. Licenses issued, prompt volumes, pilots launched and hours notionally saved all describe activity. The real question is whether AI has changed the cost structure, release cadence, capacity, revenue trajectory or organizational design of the business. Adoption measures use (and honestly, it is almost always irrelevant); value creation measures what became economically different because of that use. In the end, that means revenue growth, margin expansion, operating leverage, product velocity, new revenue streams, defensibility and enterprise value.

Productivity is only the first-order effect. Someone who works 40% faster does not improve EBITDA if the company keeps producing the same output with the same people and slightly less friction. The gain has to be converted, which we think of as a chain running from productivity to capacity to value capture. The capacity AI creates can remove cost, increase throughput, accelerate growth or improve quality, but only if management changes resource allocation, workflows or organizational structure to use it. Without those decisions, a productivity gain rarely becomes an economic one. And yes: it is harder than just hiring an AI Agency; this is where ambition kicks in.

What the company sells

The same questions extend to what the company sells. In tech for sure, but also in services and information-heavy businesses, the boundary between software and services is becoming less stable. A company that used to sell software to help customers perform a task can increasingly perform the task itself and charge for the result. A services firm can automate enough of its delivery to take on software-like economics. Work that once needed large teams can be rebuilt as software and agents, with a much smaller layer of human judgement on top. That changes the unit of economic production: for parts of the business, labor stops being the main way to scale, software and agents take over the repeatable work, and scarce human judgement moves to where it matters most. Sequoia estimates that businesses spend around six dollars on services for every dollar they spend on software, which gives a sense of how much budget comes into play once outcomes become the product. At that point the conversation moves past productivity, toward which part of the customer's work the company owns, where it sits in the value chain, and what it ultimately gets paid for.

Re-underwriting the asset

For investors, the consequences reach the investment thesis itself. AI can change the underwriting case in several ways:

- A labor-intensive operating model can become structurally leaner.
- A services business can take on software-like economics, and a software business can move from selling tools to selling outcomes.
- A product can be commoditized by a capability that did not exist at entry.
- A moat can deepen or erode, and positions in the value chain that were previously out of reach can become viable.

The same discipline applies one level down. Material AI initiatives should be underwritten like capital-allocation decisions, with a clear view of the expected economic outcome, what has to change operationally for it to appear, how quickly it should become visible, and who owns

capturing it. Starting from the P&L and working back to the technology tends to produce better choices than the reverse.

A five-year plan written eighteen months ago may already rest on assumptions that no longer hold. We therefore think portfolio companies should be re-underwritten through an AI lens on a regular basis. Individual model releases rarely change a thesis, but the frontier is moving fast enough that transformation cannot be run as a one-off project.

There is a further consequence that matters specifically to owners: AI changes the boundaries of an asset as well as its productivity. The established approach is to acquire a company and optimize the company that exists. AI adds a different set of questions:

- What should the company bring in-house that it currently buys?
- What should it automate that it currently staffs?
- What should it sell as an outcome that it currently sells as a tool?
- Which adjacent profit pools can it now enter economically?

This is closer to industrial redesign than to digital transformation. The asset itself can be re-engineered at the level of its production function, its organization and its value proposition.

Taken together, this suggests another way to assess a company, which we call its **AI transformation surface**: the economic and strategic value that could be unlocked if legacy assumptions about labor, software, workflows and product design were removed. Most AI diligence today is defensive and asks how exposed a business is. That question matters, but for an owner able to execute the transformation, a company that is technologically behind can still be an attractive asset if its customer relationships, distribution, data, domain position or workflow ownership remain valuable. Its lag is then unrealized potential as much as risk, and the useful question becomes how much value a capable owner could create by closing the gap.

How the work gets done

Not every attractive opportunity should come first. We assess opportunities on three dimensions: **value**, meaning economic materiality; **feasibility**, meaning whether it can be built with the technology, data and systems available; and **transformability**, meaning whether the organization can absorb the change. A technically elegant use case with little financial consequence is a distraction, and a large opportunity the organization cannot operationalize is unlikely to produce value either.

We also think diagnosis and execution should overlap rather than run as separate phases. Building early is often the fastest way to learn what is technically and organizationally possible. Engineering work exposes dependencies, data limitations, legacy constraints and

resistance that workshops tend to miss, and it often reveals that the real problem differs from the one first described. The work therefore moves in repeated cycles of diagnosing, building, measuring and redesigning.

Redesigning a company cannot mean putting the existing one at risk. In many cases the safer route is to build an AI-native production system alongside the current one: small, highly capable teams with a different toolchain, operating cadence and cost structure, taking on more responsibility as the model proves itself. The duplication is temporary, and the aim is a controlled migration from the old production system to the new one.

Each cycle should also make the next one cheaper and faster. Reusable components, clean interfaces, better data, evaluation systems, governance, accumulated operational context and a workforce used to working with AI all compound. This is the deeper reason to treat transformation as continuous: the lasting advantage is an organization whose capacity to change improves with every cycle.

Our perspective

At The AI Studio, we approach this as engineers. We extrapolate the playbook that is a natural fit for software orgs, where the impact of AI is already hard to ignore, but the reasoning carries well beyond software companies. As more of the enterprise runs through software, **the line between technology transformation and business transformation becomes harder to draw.**

For owners, the relevant question is shifting from whether portfolio companies are using AI to whether they are being redesigned around what AI now makes economically and technically possible.

Value creation is no longer only about optimizing the company you bought. It may increasingly be about engineering the company that should exist instead.

From thesis to execution

The opportunity is not to add more AI tools to the existing company. It is to identify where AI changes the economics of the asset, decide what should be redesigned, and build the new operating model fast enough to capture the value.

At The AI Studio, we work with private equity firms and portfolio companies at the intersection of investment thesis, operating model and engineering. We typically help in three ways:

01 Re-underwrite the asset through an AI lens

Identify where AI can materially change cost structure, capacity, product velocity, revenue potential, organizational design or the company's position in the value chain.

02 Define the transformation agenda

Separate incremental use cases from structural opportunities, prioritize them by value, feasibility and transformability, and translate them into a concrete operating and engineering roadmap.

03 Build and activate

Work alongside management and operating teams to prototype, engineer and deploy the new capabilities — using real implementation to test assumptions, expose constraints and accelerate the transformation.

Our work can begin at the fund level, during diligence, immediately post-acquisition, or inside an existing portfolio company where management wants to move beyond pilots and productivity programs.

Let's explore the AI transformation surface of your portfolio.